53-1002637-02
24 June 2013

FastIron Ethernet Switch

Administration Guide

Supporting FastIron Software Release 08.0.00a

BROCADE

Brocade Communications Systems, Incorporated

Corporate and Latin American Headquarters
Brocade Communications Systems, Inc.
130 Holger Way
San Jose, CA 95134
Tel: 1-408-333-8000
Fax: 1-408-333-8101
E-mail: info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems China HK, Ltd.
No. 1 Guanghua Road
Chao Yang District
Units 2718 and 2818
Beijing 100020, China
Tel: +8610 6588 8888
Fax: +8610 6588 9999
E-mail: china-info@brocade.com

European Headquarters
Brocade Communications Switzerland Sàrl
Centre Swissair
Tour B - 4ème étage
29, Route de l'Aéroport
Case Postale 105
CH-1215 Genève 15
Switzerland
Tel: +41 22 799 5640
Fax: +41 22 799 5641
E-mail: emea-info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems Co., Ltd. (Shenzhen WFOE)
Citic Plaza
No. 233 Tian He Road North
Unit 1308 – 13th Floor
Guangzhou, China
Tel: +8620 3891 2000
Fax: +8620 3891 2111
E-mail: china-info@brocade.com

Document History

Title	Publication number	Summary of changes	Date
FastIron Ethernet Switch Administration Guide	53-1002637-01	Release 07.4.00a has been updated with enhancements in Release 08.0.00	April 2013
FastIron Ethernet Switch Administration Guide	53-1002637-02	Release 08.0.00 has been updated with enhancements in Release 08.0.00a	June 2013

Contents

Chapter 4 Software-based Licensing

Chapter 11 Network Monitoring

Chapter 12 **Power over Ethernet**

Chapter 13 **System Monitoring**

Appendix A **Syslog messages**

Index

About This Document

This chapter contains the following sections:

Introduction

This guide includes procedures for configuring the software. The software procedures show how to perform tasks using the CLI. This guide also describes how to monitor Foundry products using statistics and summary screens.

Supported Hardware

This guide supports the following product families from Foundry:

- FastIron X Series devices (chassis models):
  - FastIron SX 800
  - FastIron SX 1600
- Brocade FCX Series (FCX) Stackable Switch
- Brocade ICX™ 6610 (ICX 6610) Stackable Switch
- Brocade ICX 6430 Series (ICX 6430)
- Brocade ICX 6450 Series (ICX 6450)
- Brocade TurboIron 24X Series

NOTE
The Brocade ICX 6430-C switch supports the same feature set as the Brocade ICX 6430 switch unless otherwise noted.

For information about the specific models and modules supported in a product family, refer to the hardware installation guide for that product family. "Related publications" on page xvi lists the hardware installation guides.

Unsupported features

Features that are not documented in "Related publications" on page xvi are not supported.

Audience

This document is designed for network engineers with a working knowledge of Layer 2 and Layer 3 switching and routing.

If you are using a Brocade Layer 3 switch, you should be familiar with the following protocols if applicable to your network – IP, RIP, OSPF, BGP, IGMP, PIM, and VRRP.

What's new in this document

This document includes the information from IronWare software release 08.0.00a. Table 1 lists the enhancements for FastIron release 08.0.00a.

Summary of enhancements in FastIron release 08.0.00a

Table 1 lists the enhancements for FastIron release 08.0.00a.

TABLE 1 Summary of enhancements in FastIron release 08.0.00a

Feature	Description	Described in
Port addressing summary	How to specify port addresses for FastIron platforms	"Specifying a port address" on page 39
Layer 3 hitless route purge	You can set the route purge timer per VRF instance. Configure the timer to set the duration for which the routes should be preserved after switchover.	"Layer 3 hitless route purge" on page 111
LLDP processing on 802.1x blocked port	This feature adds support for reception and transmission of Link Layer Discovery Protocol (LLDP) packets over an 802.1x blocked port. The default behavior is to drop received LLDP packets and not to transmit LLDP packets over an 802.1x disabled port.	"Configuring LLDP processing on 802.1x blocked port" on page 202
System Monitoring	System monitoring (sysmon) is a utility that runs as a background process and monitors connections and components of the device for specific errors and logs them. It has a default policy that controls the parameters that are monitored and actions to be taken if a fault is detected.	"System Monitoring" on page 317
Show Interface Enhancements	This releases enhances the show interface command by adding: • Show Interface Wide • Show Interface Wide Slot/Port	"Displaying the port name for an interface" on page 42
Network Time Protocol Version 4 (NTP)	Network Time Protocol Version 4 (NTPv4) feature synchronizes the local system clock in the router with the UTC.	"Network Time Protocol Version 4 (NTPv4)" on page 20

TABLE 1 **Summary of enhancements in FastIron release 08.0.00a**

Feature	Description	Described in
Default Log Buffer Size	This release enhances the default-log buffer size to 1000 lines.	"Syslog" on page 243
Enhancement to Port Group Naming	This release introduces the feature to provide a single name to a group of ports.	"Assigning port names" on page 41
Static IPv6 routes	This feature allows you to configure a static IPv6 route to be redistributed into a routing protocol.	"Static IPv6 route configuration" on page 151

Document conventions

This section describes text formatting conventions and important notice formats used in this document.

Text formatting

The narrative-text formatting conventions that are used are as follows:

bold text	Identifies command names
	Identifies the names of user-manipulated GUI elements
	Identifies keywords
	Identifies text to enter at the GUI or CLI
italic text	Provides emphasis
	Identifies variables
	Identifies document titles
`code` text	Identifies CLI output

For readability, command names in the narrative portions of this guide are presented in bold: for example, **show version**.

Command syntax conventions

Command syntax in this manual follows these conventions:

command and parameters	Commands and parameters are printed in bold.
[ ]	Optional parameter.
variable	Variables are printed in italics.
...	Repeat the previous element, for example "member[;member...]"
|	Choose from one of the parameters.

Notes, cautions, and danger notices

The following notices and statements are used in this manual. They are listed below in order of increasing severity of potential hazards.

NOTE
A note provides a tip, guidance or advice, emphasizes important information, or provides a reference to related information.

CAUTION
A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.

DANGER
A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Related publications

The following Brocade documents supplement the information in this guide and can be located at http://www.brocade.com/ethernetproducts.

- *Brocade FastIron, FCX, ICX, and TurboIron Diagnostic Reference*
- *Brocade FastIron SX Series Chassis Hardware Installation Guide*
- *Brocade FCX Series Hardware Installation Guide*
- *Brocade ICX 6430 and ICX 6450 Stackable Switch Hardware Installation Guide*
- *Brocade ICX 6610 Stackable Switch Hardware Installation Guide*
- *Brocade TurboIron 24X Series Configuration Guide*
- *Brocade TurboIron 24X Series Hardware Installation Guide*
- *FastIron Ethernet Switch IP Multicast Configuration Guide*
- *FastIron Ethernet Switch Layer 3 Routing Configuration Guide*
- *FastIron Ethernet Switch Platform and Layer 2 Switching Configuration Guide*
- *FastIron Ethernet Switch Security Configuration Guide*
- *FastIron Ethernet Switch Software Upgrade Guide*
- *FastIron Ethernet Switch Stacking Configuration Guide*
- *FastIron Ethernet Switch Traffic Management Guide*
- *Unified IP MIB Reference*
- *FastIron Ethernet Switch Software Upgrade Guide*
- *Brocade ICX 6430-C Switch Installation Guide*

Getting technical help

To contact Technical Support, go to http://www.brocade.com/services-support/index.page for the latest e-mail and telephone contact information.

Document feedback

Quality is our first concern at Brocade and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. Forward your feedback to:

documentation@brocade.com

Provide the title and version number of the document and as much detail as possible about your comment, including the topic heading and page number and your suggestions for improvement.

Chapter 1

Management Applications

Table 2 lists the individual Brocade FastIron switches and the management application features they support. These features are supported in the Layer 2 and Layer 3 software images.

TABLE 2 **Supported management application features**

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6430 ICX 6430-C12	ICX 6450
Management port	Yes (FSX 800 and FSX 1600 only)	Yes	Yes	Yes	Yes
industry-standard Command Line Interface (CLI), including support for: • Serial and Telnet access • SSH • Alias command • On-line help • Command completion • Scroll control • Line editing • Searching and filtering output • Special characters	Yes	Yes	Yes	Yes	Yes

This chapter contains the following sections:

NOTE
Configuration through web interface is not supported in this release. Only front panel display is supported using Web.

NOTE
08.0.00a release supports 5 incoming telnet/SSH sessions and 5 outgoing telnet/SSH sessions.

Management port overview

NOTE
The management port applies to FCX, SX 800, SX 1600, ICX 6430, and ICX 6450 devices.

The management port is an out-of-band port that customers can use to manage their devices without interfering with the in-band ports. The management port is widely used to download images and configurations, for Telnet sessions.

For FCX devices, the MAC address for the management port is derived from the base MAC address of the unit, plus the number of ports in the base module. For example, on a 48-port FCX standalone device, the base MAC address is 0000.0034.2200. The management port MAC address for this device would be 0000.0034.2200 plus 0x30, or 0000.0034.2230. The 0x30 in this case equals the 48 ports on the base module.

For SX 800 and SX 1600 devices, the MAC address for the management port is derived as if the management port is the last port on the management module where it is located. For example, on a 2 X 10G management module, the MAC address of the management port is that of the third port on that module.

How the management port works

The following rules apply to management ports:

- Only packets that are specifically addressed to the management port MAC address or the broadcast MAC address are processed by the Layer 2 switch or Layer 3 switch. All other packets are filtered out.
- No packet received on a management port is sent to any in-band ports, and no packets received on in-band ports are sent to a management port.
- A management port is not part of any VLAN
- Configuring a strict management VRF disables certain features on the management port.
- Protocols are not supported on the management port.
- Creating a management VLAN disables the management port on the device.
- For FCX and ICX devices, all features that can be configured from the global configuration mode can also be configured from the interface level of the management port. Features that are configured through the management port take effect globally, not on the management port itself.

For switches, any in-band port may be used for management purposes. A router sends Layer 3 packets using the MAC address of the port as the source MAC address.

For stacking devices, (for example, an FCX stack) each stack unit has one out-of band management port. Only the management port on the Active Controller will actively send and receive packets. If a new Active Controller is elected, the new Active Controller management port will become the active management port. In this situation, the MAC address of the old Active Controller and the MAC address of the new controller will be different.

CLI Commands for use with the management port

The following CLI commands can be used with a management port.

To display the current configuration, use the **show running-config interface management** command.

Syntax: **show running-config interface management** *num*

```
Brocade(config-if-mgmt)#ip addr 10.44.9.64/24
Brocade(config)#show running-config interface management 1
```

```
interface management 1
ip address 10.44.9.64 255.255.255.0
```

To display the current configuration, use the **show interfaces management** command.

Syntax: **show interfaces management** *num*

```
Brocade(config)#show interfaces management 1
GigEthernetmgmt1 is up, line protocol is up
Hardware is GigEthernet, address is 0000.0076.544a (bia 0000.0076.544a)
Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
Configured mdi mode AUTO, actual none
BPRU guard is disabled, ROOT protect is disabled
Link Error Dampening is Disabled
STP configured to OFF, priority is level0, MAC-learning is enabled
Flow Control is config disabled, oper enabled
Mirror disabled, Monitor disabled
Not member of any active trunks
Not member of any configured trunks
No port name
IPG MII 0 bits-time, IPG GMII 0 bits-time
IP MTU 1500 bytes
300 second input rate: 83728 bits/sec, 130 packets/sec, 0.01% utilization
300 second output rate: 24 bits/sec, 0 packets/sec, 0.00% utilization
39926 packets input, 3210077 bytes, 0 no buffer
Received 4353 broadcasts, 32503 multicasts, 370 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants
22 packets output, 1540 bytres, 0 underruns
Transmitted 0 broadcasts, 6 multicasts, 16 unicasts
0 output errors, 0 collisions
```

To display the management interface information in brief form, enter the **show interfaces brief management** command.

Syntax: **show interfaces brief management** *num*

```
Brocade#show interfaces brief management 1
Port   Link   State   Dupl  Speed  Trunk    Tag   Pri   MAC              Name
mgmt1  Up     None    Full  1G     None     No    0     0000.0076.544a
```

To display management port statistics, enter the **show statistics management** command.

Syntax: **show statistics management** *num*

```
Brocade#show statistics management 1
Port   Link   State   Dupl  Speed  Trunk    Tag   Pri   MAC              Name
mgmt1  Up     None    Full  1G     None     No    0     0000.0076.544a

Port mgmt1 Counters:
   InOctets3210941OutOctets1540
   InPkts39939OutPackets22
InBroadcastPkts4355OutbroadcastPkts0
InMultiastPkts35214OutMulticastPkts6
InUnicastPkts370OutUnicastPkts16
InBadPkts0
InFragments0
InDiscards0OutErrors0
CRC    0  Collisions0
InErrors0  LateCollisions0
```

```
InGiantPkts0
InShortPkts0
InJabber0
InFlowCtrlPkts0OutFlowCtrlPkts0
InBitsPerSec83728OutBitsPerSec24
InPktsPerSec130OutPktsPerSec0
InUtilization0.01%OutUtilization0.00%
```

To display the management interface statistics in brief form, enter the **show statistics brief management** command.

Syntax: **show statistics brief management** *num*

```
Brocade(config)#show statistics brief management 1
PortIn PacketsOut PacketsTrunkIn ErrorsOut Errors
mgmt1399462200

Total399452200
```

Logging on through the CLI

Once an IP address is assigned to a Brocade device running Layer 2 software or to an interface on the Brocade device running Layer 3 software, you can access the CLI either through the direct serial connection to the device or through a local or remote Telnet session.

You can initiate a local Telnet or SNMP or SSH connection by attaching a cable to a port and specifying the assigned management station IP address.

The commands in the CLI are organized into the following levels:

- User EXEC – Lets you display information and perform basic tasks such as pings and traceroutes.
- Privileged EXEC – Lets you use the same commands as those at the User EXEC level plus configuration commands that do not require saving the changes to the system-config file.
- CONFIG – Lets you make configuration changes to the device. To save the changes across reboots, you need to save them to the system-config file. The CONFIG level contains sub-levels for individual ports, for VLANs, for routing protocols, and other configuration areas.

NOTE
By default, any user who can open a serial or Telnet or SSH connection to the Brocade device can access all these CLI levels. To secure access, you can configure Enable passwords or local user accounts, or you can configure the device to use a RADIUS or TACACS/TACACS+ server for authentication. Refer to “Security Access” chapter in the *FastIron Ethernet Switch Security Configuration Guide*.

Online help

To display a list of available commands or command options, enter “?” or press Tab. If you have not entered part of a command at the command prompt, all the commands supported at the current CLI level are listed. If you enter part of a command, then enter “?” or press Tab, the CLI lists the options you can enter at this point in the command string.

If you enter an invalid command followed by ?, a message appears indicating the command was unrecognized. An example is given below.

```
Brocade(config)#rooter ip
Unrecognized command
```

Command completion

The CLI supports command completion, so you do not need to enter the entire name of a command or option. As long as you enter enough characters of the command or option name to avoid ambiguity with other commands or options, the CLI understands what you are typing. This feature is not available in the boot loader prompt of ICX 6430 and ICX 6450 devices.

Scroll control

By default, the CLI uses a page mode to paginate displays that are longer than the number of rows in your terminal emulation window. For example, if you display a list of all the commands at the global CONFIG level but your terminal emulation window does not have enough rows to display them all at once, the page mode stops the display and lists your choices for continuing the display. An example is given below.

```
aaa
all-client
appletalk
arp
boot
some lines omitted for brevity...

ipx
lock-address
logging
mac
--More--, next page: Space, next line:
Return key, quit: Control-c
```

The software provides the following scrolling options:

- Press the **Space bar** to display the next page (one screen at a time).
- Press the **Return** or **Enter** key to display the next line (one line at a time).
- Press **Ctrl+C** or **Ctrl+Q** to cancel the display.

Line editing commands

The CLI supports the following line editing commands. To enter a line-editing command, use the CTRL+key combination for the command by pressing and holding the CTRL key, then pressing the letter associated with the command.

TABLE 3 CLI line editing commands

Ctrl+Key combination	Description
Ctrl+A	Moves to the first character on the command line.
Ctrl+B	Moves the cursor back one character.

TABLE 3 CLI line editing commands (Continued)

Ctrl+Key combination	Description
Ctrl+C	Escapes and terminates command prompts and ongoing tasks (such as lengthy displays), and displays a fresh command prompt.
Ctrl+D	Deletes the character at the cursor.
Ctrl+E	Moves to the end of the current command line.
Ctrl+F	Moves the cursor forward one character.
Ctrl+K	Deletes all characters from the cursor to the end of the command line.
Ctrl+L; Ctrl+R	Repeats the current command line on a new line.
Ctrl+N	Enters the next command line in the history buffer.
Ctrl+P	Enters the previous command line in the history buffer.
Ctrl+U; Ctrl+X	Deletes all characters from the cursor to the beginning of the command line.
Ctrl+W	Deletes the last word you typed.
Ctrl+Z	Moves from any CONFIG level of the CLI to the Privileged EXEC level; at the Privileged EXEC level, moves to the User EXEC level.

Using stack-unit, slot number, and port number with CLI commands

Many CLI commands require users to enter port numbers as part of the command syntax, and many **show** command outputs display port numbers. The port numbers are entered and displayed in one of the following formats:

- port number only
- slot number and port number
- stack-unit, slot number, and port number

The following sections show which format is supported on which devices. The ports are labelled on the front panels of the devices.

CLI nomenclature on Chassis-based models

Chassis-based models (FSX 800 and FSX 1600) use port numbering that consists of a slot number and a port number. When you enter CLI commands on these devices, you must specify both the slot number and the port number. The slot numbers used in the FSX CLI examples apply only to Chassis devices.

Here is an example. The following commands change the CLI from the global CONFIG level to the configuration level for the first port on the device:

- FSX commands

```
Brocade(config)#interface e 1/1
Brocade(config-if-1/1)#
```

Syntax: ethernet *slotnum/portnum*

CLI nomenclature on Stackable devices

Stackable devices (FCX and ICX) use the **stack-unit/slot/port** nomenclature. When you enter CLI commands that include the port number as part of the syntax, you must use the **stack-unit/slot/port** number format. For example, the following commands change the CLI from the global CONFIG level to the configuration level for the first port on the device:

```
Brocade(config)#interface e 1/1/1
Brocade(config-if-e1000-1/1/1)#
```

Syntax: **ethernet** *stack-unit/slotnum/portnum*

Refer to "Brocade Stackable Devices" chapter in the *FastIron Ethernet Switch Stacking Configuration Guide* for more information about these devices.

Searching and filtering output from CLI commands

You can filter CLI output from **show** commands and at the --More-- prompt. You can search for individual characters, strings, or construct complex regular expressions to filter the output.

Searching and filtering output from Show commands

You can filter output from **show** commands to display lines containing a specified string, lines that do not contain a specified string, or output starting with a line containing a specified string. The search string is a regular expression consisting of a single character or string of characters. You can use special characters to construct complex regular expressions. Refer to "Using special characters in regular expressions" on page 9 for information on special characters used with regular expressions.

Displaying lines containing a specified string

The following command filters the output of the **show interface** command for port 3/11 so it displays only lines containing the word "Internet". This command can be used to display the IP address of the interface.

```
Brocade#show interface e 3/11 | include Internet
  Internet address is 10.168.1.11/24, MTU 1518 bytes, encapsulation ethernet
```

Syntax: *show-command* **| include** *regular-expression*

NOTE
The vertical bar (**|**) is part of the command.

Note that the regular expression specified as the search string is case sensitive. In the example above, a search string of "Internet" would match the line containing the IP address, but a search string of "internet" would not.

Displaying lines that do not contain a specified string

The following command filters the output of the **show who** command so it displays only lines that do not contain the word "closed". This command can be used to display open connections to the Brocade device.

```
Brocade#show who | exclude closed
Console connections:
        established
        you are connecting to this session
        2 seconds in idle
Telnet connections (inbound):
 1      established, client ip address 10.168.9.37
        27 seconds in idle
Telnet connection (outbound):
SSH connections:
```

Syntax: *show-command* | **exclude** *regular-expression*

Displaying lines starting with a specified string

The following command filters the output of the **show who** command so it displays output starting with the first line that contains the word "SSH". This command can be used to display information about SSH connections to the Brocade device.

```
Brocade#show who | begin SSH
SSH connections:
 1      established, client ip address 10.168.9.210
        7 seconds in idle
 2      closed
 3      closed
 4      closed
 5      closed
```

Syntax: *show-command* | **begin** *regular-expression*

Searching and filtering output at the --More-- prompt

The --More-- prompt displays when output extends beyond a single page. From this prompt, you can press the Space bar to display the next page, the Return or Enter key to display the next line, or Ctrl+C or Q to cancel the display. In addition, you can search and filter output from this prompt.

At the --More-- prompt, you can press the forward slash key (/) and then enter a search string. The Brocade device displays output starting from the first line that contains the search string, similar to the **begin** option for **show** commands. An example is given below.

```
--More--, next page: Space, next line: Return key, quit: Control-c
/telnet
```

The results of the search are displayed.

```
searching...
  telnet              Telnet by name or IP address
  temperature         temperature sensor commands
  terminal            display syslog
  traceroute          TraceRoute to IP node
  undebug             Disable debugging functions (see also 'debug')
  undelete            Undelete flash card files
  whois               WHOIS lookup
  write               Write running configuration to flash or terminal
```

To display lines containing only a specified search string (similar to the **include** option for **show** commands) press the plus sign key (+) at the --More-- prompt and then enter the search string.

```
--More--, next page: Space, next line: Return key, quit: Control-c
+telnet
```

The filtered results are displayed.

```
filtering...
  telnet                Telnet by name or IP address
```

To display lines that do not contain a specified search string (similar to the **exclude** option for **show** commands) press the minus sign key (-) at the --More-- prompt and then enter the search string.

```
--More--, next page: Space, next line: Return key, quit: Control-c
-telnet
```

The filtered results are displayed.

```
filtering...
  temperature           temperature sensor commands
  terminal              display syslog
  traceroute            TraceRoute to IP node
  undebug               Disable debugging functions (see also 'debug')
  undelete              Undelete flash card files
  whois                 WHOIS lookup
  write                 Write running configuration to flash or terminal
```

As with the commands for filtering output from **show** commands, the search string is a regular expression consisting of a single character or string of characters. You can use special characters to construct complex regular expressions. See the next section for information on special characters used with regular expressions.

Using special characters in regular expressions

You use a regular expression to specify a single character or multiple characters as a search string. In addition, you can include special characters that influence the way the software matches the output against the search string. These special characters are listed in the following table.

TABLE 4 Special characters for regular expressions

Character	Operation
.	The period matches on any single character, including a blank space. For example, the following regular expression matches “aaz”, “abz”, “acz”, and so on, but not just “az”: a.z
*	The asterisk matches on zero or more sequential instances of a pattern. For example, the following regular expression matches output that contains the string “abc”, followed by zero or more Xs: abcX*
+	The plus sign matches on one or more sequential instances of a pattern. For example, the following regular expression matches output that contains "de", followed by a sequence of “g”s, such as “deg”, “degg”, “deggg”, and so on: deg+

TABLE 4 Special characters for regular expressions (Continued)

Character	Operation
?	The question mark matches on zero occurrences or one occurrence of a pattern. For example, the following regular expression matches output that contains "dg" or "deg": de?g **NOTE:** Normally when you type a question mark, the CLI lists the commands or options at that CLI level that begin with the character or string you entered. However, if you enter Ctrl+V and then type a question mark, the question mark is inserted into the command line, allowing you to use it as part of a regular expression.
^	A caret (when not used within brackets) matches on the beginning of an input string. For example, the following regular expression matches output that begins with “deg”: ^deg
$	A dollar sign matches on the end of an input string. For example, the following regular expression matches output that ends with “deg”: deg$
_	An underscore matches on one or more of the following: • , (comma) • { (left curly brace) • } (right curly brace) • ((left parenthesis) •) (right parenthesis) • The beginning of the input string • The end of the input string • A blank space For example, the following regular expression matches on “100” but not on “1002”, “2100”, and so on. _100_
[ ]	Square brackets enclose a range of single-character patterns. For example, the following regular expression matches output that contains “1”, “2”, “3”, “4”, or “5”: [1-5] You can use the following expression symbols within the brackets. These symbols are allowed only inside the brackets. • ^ – The caret matches on any characters ***except*** the ones in the brackets. For example, the following regular expression matches output that does ***not*** contain “1”, “2”, “3”, “4”, or “5”: [^1-5] • - The hyphen separates the beginning and ending of a range of characters. A match occurs if any of the characters within the range is present. See the example above.
|	A vertical bar separates two alternative values or sets of values. The output can match one or the other value. For example, the following regular expression matches output that contains either “abc” or “defg”: abc|defg
( )	Parentheses allow you to create complex expressions. For example, the following complex expression matches on “abc”, “abcabc”, or “defg”, but not on “abcdefgdefg”: ((abc)+)|((defg)?)

If you want to filter for a special character instead of using the special character as described in the table above, enter “\” (backslash) in front of the character. For example, to filter on output containing an asterisk, enter the asterisk portion of the regular expression as “*”.

```
Brocade#show ip route bgp | include \*
```

Creating an alias for a CLI command

You can create *aliases* for CLI commands. An alias serves as a shorthand version of a longer CLI command. For example, you can create an alias called **shoro** for the CLI command **show ip route**. Then when you enter **shoro** at the command prompt, the **show ip route** command is issued.

To create an alias called **shoro** for the CLI command **show ip route**, enter the **alias shoro = show ip route** command.

```
Brocade(config)#alias shoro = show ip route
```

Syntax: [**no**] **alias** *alias-name* = *cli-command*

The *alias-name* must be a single word, without spaces.

After the alias is configured, entering **shoro** at either the Privileged EXEC or CONFIG levels of the CLI, issues the **show ip route** command.

To create an alias called **wrsbc** for the CLI command **copy running-config tftp 10.10.10.10 test.cfg**, enter the following command.

```
Brocade(config)#alias wrsbc = copy running-config tftp 10.10.10.10 test.cfg
```

To remove the **wrsbc** alias from the configuration, enter one of the following commands.

```
Brocade(config)#no alias wrsbc
```

or

```
Brocade(config)#unalias wrsbc
```

Syntax: **unalias** *alias-name*

The specified *alias-name* must be the name of an alias already configured on the Brocade device.

To display the aliases currently configured on the Brocade device, enter the following command at either the Privileged EXEC or CONFIG levels of the CLI.

```
Brocade#alias
          wrsbc      copy running-config tftp 10.10.10.10 test.cfg
           shoro       show ip route
```

Syntax: **alias**

Configuration notes for creating a command alias

The following configuration notes apply to this feature:

- You cannot include additional parameters with the alias at the command prompt. For example, after you create the **shoro** alias, **shoro bgp** would not be a valid command.
- If configured on the Brocade device, authentication, authorization, and accounting is performed on the actual command, not on the alias for the command.
- To save an alias definition to the startup-config file, use the **write memory** command.

Chapter

Basic Software Features 2

Table 5 lists the individual Brocade FastIron switches and the basic software features they support. These features are supported in the Layer 2 and Layer 3 software images, except where explicitly noted.

TABLE 5 Supported basic software features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6430	ICX 6450
Basic System Parameters					
System name, contact, and location	Yes	Yes	Yes	Yes	Yes
SNMP trap receiver and trap source address	Yes	Yes	Yes	Yes	Yes
Virtual routing interface statistics via SNMP	Yes	No	No	No	No
Disable Syslog messages and traps for CLI access	Yes	Yes	Yes	Yes	Yes
Cancelling an outbound Telnet session	Yes	Yes	Yes	Yes	Yes
Network Time Protocol Version 4 (NTPv4)	Yes[1]	Yes (on the router code only)	Yes	Yes	Yes
System clock	Yes	Yes	Yes	Yes	Yes
Byte-based broadcast, multicast, and unknown-unicast limits	Yes	No	No	No	No
Packet-based broadcast, multicast, and unknown-unicast limits	Yes	Yes	Yes	Yes	Yes
CLI banners	Yes	Yes	Yes	Yes	Yes
Local MAC address for Layer 2 management traffic	No	Yes	Yes	Yes	Yes
Show Interface Enhancements	Yes	Yes	Yes	Yes	Yes
Basic Port Parameters					
Port name	Yes	Yes	Yes	Yes	Yes
Enhancement to Port Group Naming	Yes	Yes	Yes	Yes	Yes
10/100/1000 port speed	Yes	Yes	Yes	Yes	Yes
Auto-negotiation	Yes	Yes	Yes	Yes	Yes
Auto-negotiation maximum port speed advertisement and down-shift	Yes	Yes	Yes	Yes	Yes
Duplex mode	Yes	Yes	Yes	Yes	Yes
Auto MDI/MDIX detection	Yes	Yes	Yes	Yes	Yes

TABLE 5 Supported basic software features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6430	ICX 6450
Port status (enable or disable)	Yes	Yes	Yes	Yes	Yes
Flow control: • Responds to flow control packets, but does not generate them	Yes	Yes	Yes	Yes	Yes
Symmetric flow control • Can transmit and receive 802.3x PAUSE frames	No	Yes	Yes	Yes	Yes
Auto-negotiation and advertisement of flow control	Yes	Yes	Yes	Yes	Yes
PHY FIFO Rx and TX Depth	No	Yes	Yes	Yes	Yes
Interpacket Gap (IPG) adjustment	Yes	Yes	Yes	Yes	Yes
CLI support for 100BaseTX and 100BaseFX	Yes	Yes	Yes	No	No
Gbps fiber negotiate mode	Yes	Yes	Yes	No	No
QoS priority	Yes	Yes	Yes	Yes	Yes
VOIP autoconfiguration and CDP	Yes	Yes	Yes	Yes	Yes
Port flap dampening	Yes	Yes	Yes	Yes	Yes
Port loop detection	Yes	Yes	Yes	Yes	Yes

1. Second and third generation modules.

This chapter contains the following sections:

Basic system parameter configuration

Brocade devices are configured at the factory with default parameters that allow you to begin using the basic features of the system immediately. However, many of the advanced features such as VLANs or routing protocols for the device must first be enabled at the system (global) level before they can be configured. If you use the Command Line Interface (CLI) to configure system parameters, you can find these system level parameters at the Global CONFIG level of the CLI.

NOTE
Before assigning or modifying any router parameters, you must assign the IP subnet (interface) addresses for each port.

NOTE
For information about configuring IP addresses, DNS resolver, DHCP assist, and other IP-related parameters, refer to “IP Configuration” chapter in the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide*

NOTE
For information about the Syslog buffer and messages, refer to Appendix A, "Syslog messages".

The procedures in this section describe how to configure the basic system parameters listed in Table 5.

Entering system administration information

You can configure a system name, contact, and location for a Brocade device and save the information locally in the configuration file for future reference. This information is not required for system operation but is suggested. When you configure a system name, the name replaces the default system name in the CLI command prompt.

The name, contact, and location each can be up to 255 alphanumeric characters.

Here is an example of how to configure a system name, system contact, and location.

```
Brocade(config)# hostname zappa
zappa(config)# snmp-server contact Support Services
zappa(config)# snmp-server location Centerville
zappa(config)# end
zappa# write memory
```

Syntax: **hostname** *string*

Syntax: **snmp-server contact** *string*

Syntax: **snmp-server location** *string*

The text strings can contain blanks. The SNMP text strings do not require quotation marks when they contain blanks but the host name does.

NOTE
The **chassis name** command does not change the CLI prompt. Instead, the command assigns an administrative ID to the device.

SNMP parameter configuration

Use the procedures in this section to perform the following configuration tasks:

- Specify a Simple Network Management Protocol (SNMP) trap receiver.
- Specify a source address and community string for all traps sent by the device.
- Change the holddown time for SNMP traps
- Disable individual SNMP traps. (All traps are enabled by default.)
- Disable traps for CLI access that is authenticated by a local user account, a RADIUS server, or a TACACS/TACACS+ server.

NOTE
To add and modify "get" (read-only) and "set" (read-write) community strings, refer to "Security Access" chapter in the *FastIron Ethernet Switch Security Configuration Guide*.

Specifying an SNMP trap receiver

You can specify a trap receiver to ensure that all SNMP traps sent by the Brocade device go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. When you specify the host, you also specify a community string. The Brocade device sends all the SNMP traps to the specified hosts and includes the specified community string. Administrators can therefore filter for traps from a Brocade device based on IP address or community string.

When you add a trap receiver, the software automatically encrypts the community string you associate with the receiver when the string is displayed by the CLI. If you want the software to show the community string in the clear, you must explicitly specify this when you add a trap receiver. In either case, the software does not encrypt the string in the SNMP traps sent to the receiver.

To specify the host to which the device sends all SNMP traps, use one of the following methods.

To add a trap receiver and encrypt the display of the community string, enter commands such as the following.

To specify an SNMP trap receiver and change the UDP port that will be used to receive traps, enter a command such as the following.

```
Brocade(config)# snmp-server host 10.2.2.2 0 mypublic port 200
Brocade(config)# write memory
```

Syntax: **snmp-server host** *ip-addr* [0 | 1] *string* [**port** *value*]

The *ip-addr* parameter specifies the IP address of the trap receiver.

The **0 | 1** parameter specifies whether you want the software to encrypt the string (**1**) or show the string in the clear (**0**). The default is **0**.

The *string* parameter specifies an SNMP community string configured on the Brocade device. The string can be a read-only string or a read-write string. The string is not used to authenticate access to the trap host but is instead a useful method for filtering traps on the host. For example, if you configure each of your Brocade devices that use the trap host to send a different community string, you can easily distinguish among the traps from different Brocade devices based on the community strings.

The command in the example above adds trap receiver 10.2.2.2 and configures the software to encrypt display of the community string. When you save the new community string to the startup-config file (using the **write memory** command), the software adds the following command to the file.

```
snmp-server host 10.2.2.2 1 encrypted-string
```

To add a trap receiver and configure the software to encrypt display of the community string in the CLI, enter commands such as the following.

```
Brocade(config)# snmp-server host 10.2.2.2 0 FastIron-12
Brocade(config)# write memory
```

The **port** *value* parameter allows you to specify which UDP port will be used by the trap receiver. This parameter allows you to configure several trap receivers in a system. With this parameter, a network management application can coexist in the same system. Brocade devices can be configured to send copies of traps to more than one network management application.

Specifying a single trap source

You can specify a single trap source to ensure that all SNMP traps sent by the Layer 3 switch use the same source IP address. For configuration details, refer to "Specifying a single source interface for specified packet types" section in the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide.*

Setting the SNMP trap holddown time

When a Brocade device starts up, the software waits for Layer 2 convergence (STP) and Layer 3 convergence (OSPF) before beginning to send SNMP traps to external SNMP servers. Until convergence occurs, the device might not be able to reach the servers, in which case the messages are lost.

By default, a Brocade device uses a one-minute holddown time to wait for the convergence to occur before starting to send SNMP traps. After the holddown time expires, the device sends the traps, including traps such as "cold start" or "warm start" that occur before the holddown time expires.

You can change the holddown time to a value from one second to ten minutes.

To change the holddown time for SNMP traps, enter a command such as the following at the global CONFIG level of the CLI.

```
Brocade(config)# snmp-server enable traps holddown-time 30
```

The command in this example changes the holddown time for SNMP traps to 30 seconds. The device waits 30 seconds to allow convergence in STP and OSPF before sending traps to the SNMP trap receiver.

Syntax: [**no**] **snmp-server enable traps holddown-time** *secs*

The *secs* parameter specifies the number of seconds and can be from 1 – 600 (ten minutes). The default is 60 seconds.

Disabling SNMP traps

Brocade devices come with SNMP trap generation enabled by default for all traps. You can selectively disable one or more of the following traps.

NOTE
By default, all SNMP traps are enabled at system startup.

SNMP Layer 2 traps

The following traps are generated on devices running Layer 2 software:

- SNMP authentication keys
- Power supply failure
- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root

- Bridge topology change
- Locked address violation

SNMP Layer 3 traps

The following traps are generated on devices running Layer 3 software:

- SNMP authentication key
- Power supply failure
- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root
- Bridge topology change
- Locked address violation
- BGP4
- OSPF
- VRRP
- VRRP-E

To stop link down occurrences from being reported, enter the following.

```
Brocade(config)# no snmp-server enable traps link-down
```

Syntax: [**no**] **snmp-server enable traps** *trap-type*

SNMP ifIndex

On Brocade IronWare devices, SNMP Management Information Base (MIB) uses Interface Index (ifIndex) to assign a unique value to each port on a module or slot. The number of indexes that can be assigned per module is 64. On all IronWare devices, the system automatically assign 64 indexes to each module on the device. This value is not configurable.

Displaying virtual routing interface statistics

NOTE
This feature is supported on FastIron X Series devices only.

You can enable SNMP to extract and display virtual routing interface statistics from the ifXTable (64-bit counters).

The following describes the limitations of this feature:

- The Brocade device counts traffic from all virtual interfaces (VEs). For example, in a configuration with two VLANs (VLAN 1 and VLAN 20) on port 1, when traffic is sent on VLAN 1, the counters (VE statistics) increase for both VE 1 and VE 20.
- The counters include all traffic on each virtual interface, even if the virtual interface is disabled.
- The counters include traffic that is denied by ACLs or MAC address filters.

To enable SNMP to display VE statistics, enter the **enable snmp ve-statistics** command.

```
Brocade(config)# enable snmp ve-statistics
```

Syntax: [**no**] **enable snmp ve-statistics**

Use the **no** form of the command to disable this feature once it is enabled.

Note that the above CLI command enables SNMP to display virtual interface statistics. It does not enable the CLI to display the statistics.

Disabling Syslog messages and traps for CLI access

Brocade devices send Syslog messages and SNMP traps when a user logs into or out of the User EXEC or Privileged EXEC level of the CLI. The feature applies to users whose access is authenticated by an authentication-method list based on a local user account, RADIUS server, or TACACS/TACACS+ server.

NOTE
The Privileged EXEC level is sometimes called the "Enable" level, because the command for accessing this level is **enable**.

The feature is enabled by default.

Examples of Syslog messages for CLI access

When a user whose access is authenticated by a local user account, a RADIUS server, or a TACACS or TACACS+ server logs into or out of the CLI User EXEC or Privileged EXEC mode, the software generates a Syslog message and trap containing the following information:

- The time stamp
- The user name
- Whether the user logged in or out
- The CLI level the user logged into or out of (User EXEC or Privileged EXEC level)

NOTE
Messages for accessing the User EXEC level apply only to access through Telnet. The device does not authenticate initial access through serial connections but does authenticate serial access to the Privileged EXEC level. Messages for accessing the Privileged EXEC level apply to access through the serial connection or Telnet.

The following examples show login and logout messages for the User EXEC and Privileged EXEC levels of the CLI.

```
Brocade# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
Buffer logging: level ACDMEINW, 12 messages logged
level code: A=alert C=critical D=debugging M=emergency E=error
I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed

Dynamic Log Buffer (50 entries):
Oct 15 18:01:11:info:dg logout from USER EXEC mode
Oct 15 17:59:22:info:dg logout from PRIVILEGE EXEC mode
Oct 15 17:38:07:info:dg login to PRIVILEGE EXEC mode
Oct 15 17:38:03:info:dg login to USER EXEC mode
```

Syntax: **show logging**

The first message (the one on the bottom) indicates that user "dg" logged in to the CLI User EXEC level on October 15 at 5:38 PM and 3 seconds (Oct 15 17:38:03). The same user logged into the Privileged EXEC level four seconds later.

The user remained in the Privileged EXEC mode until 5:59 PM and 22 seconds. (The user could have used the CONFIG modes as well. Once you access the Privileged EXEC level, no further authentication is required to access the CONFIG levels.) At 6:01 PM and 11 seconds, the user ended the CLI session.

Disabling the Syslog messages and traps

Logging of CLI access is enabled by default. If you want to disable the logging, enter the following commands.

```
Brocade(config)# no logging enable user-login
Brocade(config)# write memory
Brocade(config)# end
Brocade# reload
```

Syntax: [**no**] **logging enable user-login**

Cancelling an outbound Telnet session

If you want to cancel a Telnet session from the console to a remote Telnet server (for example, if the connection is frozen), you can terminate the Telnet session by doing the following.

1. At the console, press **Ctrl+^** (Ctrl+Shift-6).
2. Press the **X** key to terminate the Telnet session.

Pressing **Ctrl+^** twice in a row causes a single **Ctrl+^** character to be sent to the Telnet server. After you press **Ctrl+^**, pressing any key other than **X** or **Ctrl+^** returns you to the Telnet session.

Network Time Protocol Version 4 (NTPv4)

NTPv4 feature synchronizes the local system clock in the device with the UTC. The synchronization is achieved by maintaining a loop-free timing topology computed as a shortest-path spanning tree rooted on the primary server. NTP does not know about local time zones or daylight-saving time. A time server located anywhere in the world can provide synchronization to a client located anywhere

else in the world. It allows clients to use different time zone and daylight-saving properties. Primary servers are synchronized by wire or radio to national standards such as GPS. Timing information is conveyed from primary servers to secondary servers and clients in the network. NTP runs on UDP, which in turn runs on IP.

NTP has a hierarchical structure. NTP uses the concept of a stratum to describe how many NTP hops away a machine is from an authoritative time source. A stratum 1 time server typically has an authoritative time source such as a radio or atomic clock, or a Global Positioning System [GPS] time source directly attached. A stratum 2 time server receives its time through NTP from a stratum 1 time server and so on. As the network introduces timing discrepancies, lower stratum devices are a factor less accurate. A hierarchical structure allows the overhead of providing time to many clients to be shared among many time servers. Not all clients need to obtain time directly from a stratum 1 reference, but can use stratum 2 or 3 references.

NTP operates on a client-server basis. The current implementation runs NTP as a secondary server and/or a NTP Client. As a secondary server, the device operates with one or more upstream servers and one or more downstream servers or clients. A client device synchronizes to one or more upstream servers, but does not provide synchronization to dependant clients. Secondary servers at each lower level are assigned stratum numbers one greater than the preceding level. As stratum number increases, the accuracy decreases. Stratum one is assigned to Primary servers.

NTP uses the concept of associations to describe communication between two machines running NTP. NTP associations are statistically configured. On startup or on the arrival of NTP packets, associations are created. Multiple associations are created by the protocol to communicate with multiple servers. NTP maintains a set of statistics for each of the server or the client it is associated with. The statistics represent measurements of the system clock relative to each server clock separately. NTP then determines the most accurate and reliable candidates to synchronize the system clock. The final clock offset applied for clock adjustment is a statistical average derived from the set of accurate sources.

When multiple sources of time (hardware clock, manual configuration) are available, NTP is always considered to be more authoritative. NTP time overrides the time that is set by any other method.

NTPv4 obsoletes NTPv3 (RFC1305) and SNTP (RFC4330). SNTP is a subset of NTPv4. RFC 5905 describes NTPv4.

To keep the time in your network current, it is recommended that each device have its time synchronized with at least four external NTP servers. External NTP servers should be synchronized among themselves to maintain time synchronization.

NOTE
Network Time Protocol (NTP) commands must be configured on each individual device.

FIGURE 1 NTP Hierarchy

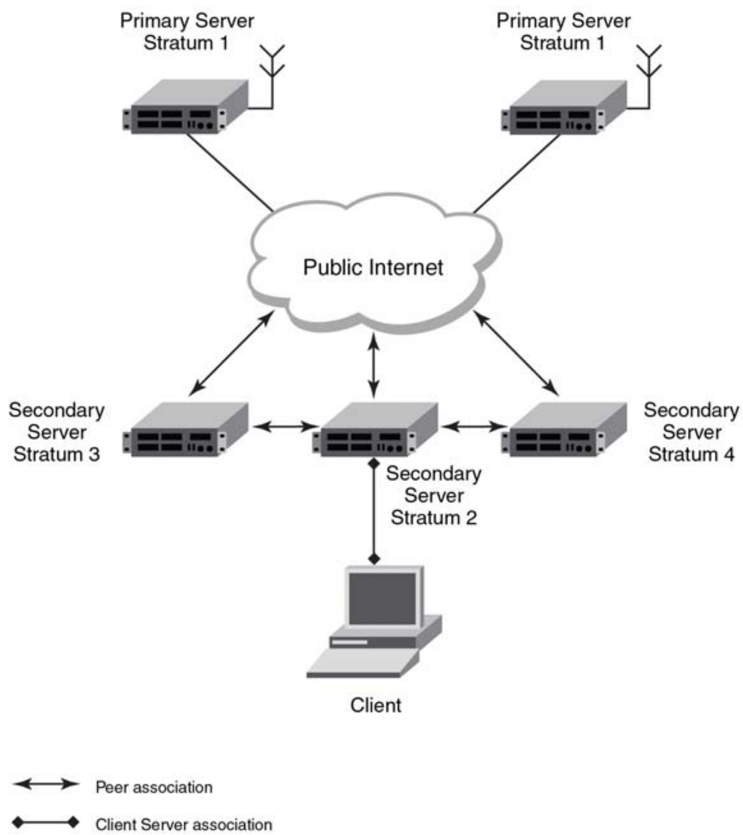

- NTP implementation conforms to RFC 5905.
- NTP can be enabled in server and client mode simultaneously.
- The NTP uses UDP port 123 for communicating with NTP servers/peers.
- NTP server and client can communicate using IPv4 or IPv6 address
- NTP implementation supports below association modes.
  - Client
  - Server
  - Symmetric active/passive
  - Broadcast server

  - Broadcast client
- NTP supports maximum of 8 servers and 8 peers. The 8 peers includes statically configured and dynamically learned.
- NTP can operate in authenticate or non-authenticate mode. Only symmetric key authentication is supported.
- By default, NTP operates in default VLAN and it can be changed.

Limitations

- FastIron devices cannot operate as primary time server (or stratum 1). It only serves as secondary time server (stratum 2 to 15).
- NTP server and client cannot communicate using hostnames.
- NTP is not supported on VRF enabled interface or ve.
- Autokey public key authentication is not supported.
- The NTP version 4 Extension fields are not supported. The packets containing the extension fields are discarded.
- The NTP packets having control (6) or private (7) packet mode is not supported. NTP packets with control and private modes will be discarded.
- On reboot or switchover, all the NTP state information will be lost and time synchronization will start fresh.
- NTP multicast server/client and manycast functionalities are not supported.
- NTP versions 1 and 2 are not supported.
- NTP MIB is not supported.

NTP and SNTP

FastIron 07.3.00c and earlier releases implements SNTP for time synchronization. In FastIron 07.3.00d, NTP can be used for time synchronization in FCX devices with router images. From FastIron 8.0 release onwards, NTP can be used for time synchronization in all FastIron devices with both router and switch images.

NTP and SNTP implementations cannot operate at the same time and one of them has to be disabled.

On downgrading from FastIron 07.3.00d to FastIron 07.3.00c or lower version, the entire NTP configuration is lost.

NTP server

A NTP server will provide the correct network time on your device using the Network time protocol (NTP). Network Time Protocol can be used to synchronize the time on devices across a network. A NTP time server is used to obtain the correct time from a time source and adjust the local time in each connecting device.

The NTP server functionality is enabled when you use the ntp command, provided SNTP configuration is already removed.

When the NTP server is enabled, it will start listening on the NTP port for client requests and responds with the reference time. Its stratum number will be the upstream time server's stratum + 1. The stratum 1 NTP server is the time server which is directly attached to the authoritative time source.

The device cannot be configured as primary time server with stratum 1. It can be configured as secondary time server with stratum 2 to 15 to serve the time using the local clock.

The NTP server is stateless and will not maintain any NTP client information.

System as an Authoritative NTP Server

The NTP server can operate in master mode to serve time using the local clock, when it has lost synchronization. Serving local clock can be enabled using the master command. In this mode, the NTP server stratum number is set to the configured stratum number. When the master command is configured and the device was never synchronized with an upstream time server and the clock setting is invalid, the server will respond to client's request with the stratum number set to 16. While the device is operating in the master mode and serving the local clock as the reference time, if synchronization with the upstream server takes place it will calibrate the local clock using the NTP time. The stratum number will switch to that of the synchronized source +1. And when synchronization is lost, the device switches back to local clock time with stratum number as specified manually (or the default).

NOTE
Local time and time zone has to be configured before configuring the master command.

- The following scenarios are observed when the master command is not configured and the NTP upstream servers are configured:
- If the synchronization with the NTP server/peer is active, the system clock is synchronized and the reference time is the NTP time.
- If the NTP server/peer is configured but not reachable and if the local clock is valid, the server will respond to client's request with the stratum number set to 16.
- If there is no NTP server/peer configured and if the local clock is valid, the server will respond to client's request with the stratum number set to 16.
- If there is no NTP server/peer configured and if the local clock is invalid, the system clock is not synchronized.

The following scenarios are observed when the master command is configured and the NTP upstream servers are also configured:

- If the synchronization with the time server/peer is active, system clock is synchronized and the reference time is the NTP time.If the NTP server/peer is configured but not reachable, the system clock is synchronized. If the local time is valid then the reference time is the local clock time.
- If the NTP server/peer is not configured, the system clock is synchronized. If the local clock is valid, then the reference time is the local clock time.
- If the NTP server/peer is not configured and the local clock is invalid, system clock is not synchronized.

NOTE
Use the master command with caution. It is very easy to override valid time sources using this command, especially if a low stratum number is configured. Configuring multiple machines in the same network with the master command can cause instability in timekeeping if the machines do not agree on the time.

NTP Client

An NTP client gets time responses from an NTP server or servers, and uses the information to calibrate its clock. This consists of the client determining how far its clock is off and adjusting its time to match that of the server. The maximum error is determined based on the round-trip time for the packet to be received.

The NTP client can be enabled when we enter the **ntp** command and configure one or more NTP servers/peers.

The NTP client maintains the server and peer state information as association. The server and peer association is mobilized at the startup or whenever user configures. The statically configured server/peer associations are not demobilized unless user removes the configuration. The symmetric passive association is mobilized upon arrival of NTP packet from peer which is not statically configured. The associations will be demobilized on error or time-out.

NTP peer

NTP peer mode is intended for configurations where a group of devices operate as mutual backups for each other. If one of the devices loses a reference source, the time values can flow from the surviving peers to all the others. Each device operates with one or more primary reference sources, such as a radio clock, or a subset of reliable NTP secondary servers. When one of the devices lose all reference sources or simply cease operation, the other peers automatically reconfigures so that time values can flow from the surviving peers to others.

When the NTP server or peer is configured with burst mode, client will send burst of up to 8 NTP packets in each polling interval. The burst number of packets in each interval increases as the polling interval increases from minimum polling interval towards maximum interval.

The NTP peer can operate in:

- Symmetric Active-When the peer is configured using the peer command.
- Symmetric Passive-Dynamically learned upon arrival of a NTP packet from the peer which is not configured. The symmetric passive association is removed on timeout or error.

The following scenarios are observed when the upstream server is not reachable after retries:

- If the NTP server/peer is configured and the master command is not configured, then the system clock is synchronized. When the system clock is synchronized, the server will respond to client's request with the stratum number set to +1. And when the system clock is unsynchronized, the server will respond to client's request with the stratum number set to 16.
- If the NTP server/peer is configured and the master command is configured, then the system clock is synchronized. When the system clock is synchronized, the reference time is the local clock time. If the local clock is valid then the server will respond to client's request with the specified stratum number if it is configured otherwise with the default stratum number.

The following scenarios are observed when you remove the last NTP server/peer under the conditions - the NTP server/peer is configured, master command is not configured, system clock is synchronized and the reference time is the NTP time:

- If the local clock is not valid, the system clock is not synchronized.
- If the local clock is valid, the system clock is synchronized and the reference time is the local clock. The server will respond to the client's request with the specified stratum number if it is configured otherwise with the default stratum number.

NOTE
To create a symmetric active association when a passive association is already formed, disable NTP, configure peer association and then enable NTP again.

NTP broadcast server

An NTP server can also operate in a broadcast mode. Broadcast servers send periodic time updates to a broadcast address, while multicast servers send periodic updates to a multicast address. Using broadcast packets can greatly reduce the NTP traffic on a network, especially for a network with many NTP clients.

The interfaces should be enabled with NTP broadcasting. The NTP broadcast server broadcasts the

NTP packets periodically (every 64 sec) to subnet broadcast IP address of the configured interface.

- NTP broadcast packets are sent to the configured subnet when the NTP broadcast server is configured on the interface which is up and the IP address is configured for the broadcast subnet under the following conditions:
  - The local clock is valid and the system clock is synchronized
  - The local clock is valid and the system clock is not synchronized
  - Authentication key is configured, the system clock is synchronized and the local clock is valid
- NTP broadcast packets are not sent in the following cases:
  - NTP broadcast server is configured on the interface which is down even if the system clock is synchronized and the local clock is valid.
  - NTP broadcast server is configured on the interface which is up and no IP address is configured for the broadcast subnet even if the system clock is synchronized and the local clock is valid.
  - NTP broadcast server is configured on the interface which is not present and no IP address is configured for the broadcast subnet even if the system clock is synchronized and the local clock is valid.
  - NTP broadcast server without authentication key is configured on the interface which is up and the IP address is configured for the broadcast subnet even when NTP authentication is enforced and the system clock is synchronized and the local clock is valid.

NTP broadcast client

An NTP broadcast client listens for NTP packets on a broadcast address. When the first packet is received, the client attempts to quantify the delay to the server, to better quantify the correct time from later broadcasts. This is accomplished by a series of brief interchanges where the client and server act as a regular (non-broadcast) NTP client and server. Once interchanges occur, the client has an idea of the network delay and thereafter can estimate the time based only on broadcast packets.

NTP associations

Networking devices running NTP can be configured to operate in variety of association modes when synchronizing time with reference time sources. A networking device can obtain time information on a network in two ways-by polling host servers and by listening to NTP broadcasts. That is, there are two types of associations-poll-based and broadcast-based.

NTP poll-based associations

The following modes are the NTP polling based associations:

1. Server mode
2. Client mode
3. Symmetric Active/Passive

The server mode requires no prior client configuration. The server responds to client mode NTP packets. Use the master command to set the device to operate in server mode when it has lost the synchronization.

When the system is operating in the client mode, it polls all configured NTP servers and peers. The device selects a host from all the polled NTP servers to synchronize with. Because the relationship that is established in this case is a client-host relationship, the host will not capture or use any time information sent by the local client device. This mode is most suited for file-server and workstation clients that are not required to provide any form of time synchronization to other local clients. Use the server and peer to individually specify the time server that you want the networking device to consider synchronizing with and to set your networking device to operate in the client mode.

Symmetric active/passive mode is intended for configurations where group devices operate as mutual backups for each other. Each device operates with one or more primary reference sources, such as a radio clock, or a subset of reliable NTP secondary servers. If one of the devices lose all reference sources or simply cease operation, the other peers automatically reconfigures. This helps the flow of time value from the surviving peers to all the others.

When a networking device is operating in the symmetric active mode, it polls its assigned time-serving hosts for the current time and it responds to polls by its hosts. Because symmetric active mode is a peer-to-peer relationship, the host will also retain time-related information of the local networking device that it is communicating with. When many mutually redundant servers are interconnected via diverse network paths, the symmetric active mode should be used. Most stratum 1 and stratum 2 servers on the Internet adopt the symmetric active form of network setup. The FastIron device operates in symmetric active mode, when the peer information is configured using the peer command and specifying the address of the peer. The peer is also configured in symmetric active mode in this way by specifying the FastIron device information. If the peer is not specifically configured, a symmetric passive association is activated upon arrival of a symmetric active message.

The specific mode that you should set for each of your networking devices depends primarily on the role that you want them to assume as a timekeeping device (server or client) and the device's proximity to a stratum 1 timekeeping server. A networking device engages in polling when it is operating as a client or a host in the client mode or when it is acting as a peer in the symmetric active mode. An exceedingly large number of ongoing and simultaneous polls on a system can seriously impact the performance of a system or slow the performance of a given network. To avoid having an excessive number of ongoing polls on a network, you should limit the number of direct, peer-to-peer or client-to-server associations. Instead, you should consider using NTP broadcasts to propagate time information within a localized network.

NTP broadcast-based associations

The broadcast-based NTP associations should be used in configurations involving potentially large client population. Broadcast-based NTP associations are also recommended for use on networks that have limited bandwidth, system memory, or CPU resources.

The devices operating in the broadcast server mode broadcasts the NTP packets periodically which can be picked up by the devices operating in broadcast client mode. The broadcast server is configured using the **broadcast** command.

A networking device operating in the broadcast client mode does not engage in any polling. Instead, the device receives the NTP broadcast server packets from the NTP broadcast servers in the same subnet. The NTP broadcast client forms a temporary client association with the NTP broadcast server. A broadcast client is configured using the **broadcast client** command. For broadcast client mode to work, the broadcast server and the clients must be located on the same subnet.

Synchronizing time

After the system peer is chosen, the system time is synchronized based on the time difference with system peer:

- If the time difference with the system peer is 128 msec and < 1000 sec, the system clock is stepped to the system peer reference time and the NTP state information is cleared.

Authentication

The time kept on a machine is a critical resource, so it is highly recommended to use the encrypted authentication mechanism.

The NTP can be configured to provide cryptographic authentication of messages with the clients/peers, and with its upstream time server. Symmetric key scheme is supported for authentication. The scheme uses MD5 keyed hash algorithm.

The authentication can be enabled using the **authenticate** command. The set of symmetric key and key string is specified using the **authentication-key** command.

If authentication is enabled, NTP packets not having a valid MAC address are dropped.

If the NTP server/peer is configured without authentication keys, the NTP request is not sent to the configured server/peer.

NOTE
The same set or subset of key id and key string should be installed on all NTP devices.

VLAN and NTP

When VLAN is configured,

- NTP time servers should be reachable through the interfaces which belong to the configured VLAN. Otherwise, NTP packets are not transmitted. This is applicable to both the unicast and the broadcast server/client.
- NTP broadcast packets are sent only on the interface which belongs to the configured VLAN.
- The received unicast or broadcast NTP packet are dropped if the interface on which packet has been received does not belong to the configured VLAN

Configuring NTP

NTP services are disabled on all interfaces by default.

Prerequisites:

- Before you begin to configure NTP, you must use the clock set command to set the time on your device to within 1000 seconds of the coordinated Universal Time (UTC).
- Disable SNTP by removing all the SNTP configurations.

Enabling NTP

NTP and SNTP implementations cannot operate simultaneously. By default, SNTP is enabled. To disable SNTP and enable NTP, use the **ntp** command in configuration mode. This command enables the NTP client and server mode if SNTP is disabled.

```
Brocade(config)# ntp
Brocade(config-ntp)#
```

Syntax: Syntax: [no] ntp

Use the **no** form of the command to disable NTP and remove the NTP configuration.

NOTE
The **no ntp** command removes all the configuration which are configured statistically and learned associations from NTP neighbors.

NOTE
You cannot configure the ntp command if SNTP is enabled. If SNTP is enabled, configuring the ntp command will display the following message:`"SNTP is enabled. Disable SNTP before using NTP for time synchronization"`

Disabling NTP

To disable the NTP server and client mode, use the **disable** command in NTP configuration mode. Disabling the NTP server or client mode will not remove the configurations.

```
Brocade(config-ntp)# disable
```

Syntax: [no] disable [serve]

If the serve keyword is specified, then NTP will not serve the time to downstream devices. The *serve* keyword disables the NTP server mode functionalities. If the serve keyword is not specified, then both NTP client mode and NTP server mode functionalities are disabled.

Use the no form of the command to enable NTP client and server mode. To enable the client mode, use the **no disable** command. To enable the client and server mode, use the **no disable serve** command. The **no disable** command enables both client and server, if the client is already enabled and server is disabled at that time “no disable server ” enables the server.

NOTE
The **disable** command disables the NTP server and client mode; it does not remove the NTP configuration.

Enabling NTP authentication

To enable Network Time Protocol (NTP) strict authentication, use the **authenticate** command. To disable the function, use the **no** form of this command.

By default, authentication is disabled.

```
Brocade(config-ntp)# [no] authenticate
```

Syntax: [no] authenticate

Defining an authentication key

To define an authentication key for Network Time Protocol (NTP), use the **authentication-key** command. To remove the authentication key for NTP, use the **no** form of this command.

By default, authentication keys are not configured.

```
Brocade(config-ntp)# authentication-key key-id 1 md5 moof
```

Syntax: [no] authentication-key *key-id* **md5** *key-string*

The valid key-id parameter is 1 to 65535.

MD5 is the message authentication support that is provided using the Message Digest 5 Algorithm. The key type md5 is currently the only key type supported.

The key-string option is the value of the MD5 key. The maximum length of the key string may be defined up to 16 characters. Up to 32 keys may be defined.

Specifying a source interface

When the system sends an NTP packet, the source IP address is normally set to the address of the interface through which the NTP packet is sent. Use the **source-interface** command to configure a specific interface from which the IP source address will be taken. To remove the specified source address, use the no form of this command.

This interface will be used for the source address for all packets sent to all destinations. If a source address is to be used for a specific association, use the source keyword in the peer or server command.

NOTE
If the source-interface is not configured, then the lowest IP address in the outgoing interface will be used in the NTP packets. Source IP address of a tunnel interface is not supported.

```
Brocade(config-ntp)# source-interface ethernet 1/3/1
```

Syntax: [no] source-interface ethernet *port* **| loopback** *num* | **ve** *num*

Specify the *port* parameter in the format *stack-unit/slotnum/portnum.*

The **loopback** *num* parameter specifies the loopback interface number.

The **ve** num parameter specifies the virtual port number.

Enable or disable the VLAN containment for NTP

To enable or disable the VLAN containment for NTP, use the **access-control vlan** command. To remove the specified NTP VLAN configuration, use the no form of this command.

NOTE
The management interface is not part of any VLAN. When configuring the VLAN containment for NTP, it will not use the management interface to send or receive the NTP packets.

```
Brocade(config-ntp)# access-control vlan 100
```

Syntax: [no] access-control vlan *vlan-id*

The *vlan-id* parameter specifies the VLAN ID number.

Configuring the NTP client

To configure the device in client mode and specify the NTP servers to synchronize the system clock, use the server command. A maximum 8 NTP servers can be configured. To remove the NTP server configuration, use the **no** form of this command.

By default, no servers are configured.

```
Brocade(config-ntp)#server 1.2.3.4 key 1234
```

Syntax: [no] server *ipv4-address* **|** *ipv6-address* [**version** *num]* [**key** *key-id]* [**minpoll** *interval*] [**maxpoll** *interval*] [**burst**]

The *ipv4-address or ipv6-address* parameter is the IP address of the server providing the clock synchronization.

The **version** *num* option defines the Network Time Protocol (NTP) version number. Valid values are 3 or 4. If the num option is not specified, the default is 4.

The **key** *key-id* option defines the authentication key. By default, no authentication key is configured.

The **minpoll** *interval* option is the shortest polling interval. The range is from 4 through 17. Default is 6. The interval argument is power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s, and so on).

The **maxpoll** *interval* option is the longest polling interval. The range is 4 through 17. Default is 10. The interval argument is calculated by the power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s, and so on).

The **burst** option sends a burst of packets to the server at each polling interval.

Configuring the master

To configure the FastIron device as a Network Time Protocol (NTP) master clock to which peers synchronize themselves when an external NTP source is not available, use the **master** command. The master clock is disabled by default. To disable the master clock function, use the no form of this command.

NOTE
This command is not effective, if the NTP is enabled in client-only mode.

```
Brocade(config-ntp)# master stratum 5
```

Syntax: [**no**] **master** [**stratum** *number*]

The number variable is a number from 2 to 15. It indicates the NTP stratum number that the system will claim.

Configuring the NTP peer

To configure the software clock to synchronize a peer or to be synchronized by a peer, use the peer command. A maximum of 8 NTP peers can be configured. To disable this capability, use the **no** form of this command.

This **peer** command is not effective if the NTP is enabled in client-only mode.

NOTE
If the peer is a member of symmetric passive association, then configuring the **peer** command will fail.

```
Brocade(config-ntp)# peer 1.2.3.4 key 1234
```

Syntax: [**no**] **peer** *ipv4-address* | *ipv6-address* [**version** *num* [**key** *key-id]* [**minpoll** *interval]* [**maxpoll** *interval]* [**burst**]

The *ipv4-address or ipv6-address* parameter is the IP address of the peer providing the clock synchronization.

The **version** *num* option defines the Network Time Protocol (NTP) version number. Valid values are 3 and 4. If this option is not specified, then the default is 4.

The **key** *key-id* option defines the authentication key. By default, no authentication key is configured.

The **minpoll** *interval* option is the shortest polling interval. The range is from 4 through 17. Default is 6. The interval argument is power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s, and so on).

The **maxpoll** *interval* option is the longest polling interval. The range is 4 through 17. Default is 10. The interval argument is calculated by the power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s, and so on).

The **burst** option sends a burst of packets to the peer at each polling interval.

NOTE
When the NTP server/peer is configured, the **master** command is not configured; on configuring the **clock set** command the system clock is not synchronized. When the **master** command is configured, on configuring the **clock set** command the system clock is synchronized and the reference time will be the local clock.

To have active peers at both the ends, you need to disable NTP, configure the peers and enable the NTP using the **no disable** command.

Configuring NTP on an interface

To configure the NTP interface context, use the **ntp-interface** command. The broadcast server or client is configured on selected interfaces. To remove the NTP broadcast configurations on the specified interface, use the no form of this command.

NOTE
The **ntp-interface** command is a mode change command, and will not be included in to the show run output unless there is configuration below that interface.

```
Brocade(config-ntp)# ntp-interface ethernet 2/13
Brocade(config-ntp-if-e1000-2/13)# exit
Brocade(config-ntp)# ntp-interface management 1
Brocade(config-ntp-mgmt-1)# exit
Brocade(config-ntp)# ntp-interface ve 100
Brocade(config-ntp-ve-100)#
```

Syntax: [**no**] **ntp-interface management 1** | **ethernet** *port* | **ve** *id*

The **management 1** parameter is the management port 1.

The **ethernet** *port* parameter specifies the ethernet port number. Specify the port parameter in the format *stack-unit/slotnum/portnum.*

The **ve** *id* parameter specifies the virtual port number.

Configuring the broadcast client

To configure a device to receive Network Time Protocol (NTP) broadcast messages on a specified interface, use the broadcast client command. NTP broadcast client can be enabled on maximum of 16 ethernet interfaces. If the interface is operationally down or NTP is disabled, then the NTP broadcast server packets are not received. To disable this capability, use the **no** form of this command.

```
Brocade(config-ntp mgmt-1)# broadcast client
```

Syntax: [**no**] **broadcast client**

Configuring the broadcast destination

To configure the options for broadcasting Network Time Protocol (NTP) traffic, use the **ntp broadcast destination** command. The NTP broadcast server can be enabled on maximum 16 ethernet interfaces and four subnet addresses per interface. If the interface is operationally down or there is no ip address configured for the subnet address, then the NTP broadcast server packets are not sent. To disable this capability, use the no form of this command.

By default, the broadcast mode is not enabled.

NOTE
This command is not effective, if the NTP server is disabled.

```
Brocade(config)#int m1
Brocade(config-if-mgmt-1)#ip address 10.20.99.173/24
Brocade(config-if-mgmt-1)#ntp
Brocade(config-ntp)#ntp-interface m1
Brocade(config-ntp -mgmt-1)# broadcast destination 10.20.99.0 key 2
```

Syntax: [**no**] **broadcast destination** *ip-address* [**key** *key-id]* [**version** *num*]

The *ip-address* parameter is the IPv4 subnet address of the device to send NTP broadcast messages to.

The **key** *key-id* option defines the authentication key. By default, no authentication key is configured.

The **version** *num* option defines the Network Time Protocol (NTP) version number. If this option is not specified, then the default value is 4.

Displaying NTP status

Use the **show ntp status** command to display the NTP status.

```
Brocade#show ntp status
Clock is synchronized, stratum 4, reference clock is 10.20.99.174
precision is 2**-16
reference time is D281713A.80000000 (03:21:29.3653007907 GMT+00 Thu Dec 01 2011)
clock offset is -2.3307 msec, root delay is 24.6646 msec
root dispersion is 130.3376 msec, peer dispersion is 84.3335 msec
system poll interval is 64, last clock update was 26 sec ago
NTP server mode is enabled, NTP client mode is enabled
NTP master mode is disabled, NTP master stratum is 8
NTP is not in panic mode
```

Table 6 **show ntp status** command output descriptions

TABLE 6 NTP status command output descriptions

Field	Description
synchronized	Indicates the system clock is synchronized to NTP server or peer.
stratum	Indicates the stratum number that this system is operating. Range 2..15.
reference	IPv4 address or first 32 bits of the MD5 hash of the IPv6 address of the peer to which clock is synchronized.
precision	Precision of the clock of this system in Hz.
reference time	Reference time stamp.
clock offset	Offset of clock (in milliseconds) to synchro-nized peer.
root delay	Total delay (in milliseconds) along path to root clock.
root dispersion	Dispersion of root path.

TABLE 6 NTP status command output descriptions

Field	Description
peer dispersion	Dispersion of root path.
system poll interval	Poll interval of the local system.
last update	Time the router last updated its NTP information.
server mode	Status of the NTP server mode for this device.
client mode	Status of the NTP client mode for this device.
master	Status of the master mode.
master stratum	Stratum number that will be used by this device when master is enabled and no upstream time servers are accessible.
panic mode	Status of the panic mode.

Displaying NTP associations

Use the **show ntp associations** command to display detailed association information of the NTP server or peers.

```
Brocade# show ntp associations
address ref clock st when poll reach delay offset disp
*~172.19.69.1 172.24.114.33 3 25 64 3 2.89 0.234 39377
~2001:235::234
INIT 16 - 64 0 0.00 0.000 15937
* synced, # selected, + candidate, - outlayer, x falseticker, ~ configured
```

Table 7 **show ntp associations** command output descriptions

TABLE 7 NTP associations command output descriptions

Field	Description
*	The peer has been declared the system peer and lends its variables to the system variables.
#	This peer is a survivor in the selection algorithm.
+	This peer is a candidate in the combine algorithm.
-	This peer is discarded as outlier in the clustering algorithm.
x	This peer is discarded as 'falseticker' in the selection algorithm.
~	The server or peer is statically configured.
address	IPv4 or IPv6 address of the peer.
ref clock	IPv4 address or first 32 bits of the MD5 hash of the IPv6 address of the peer to which clock is synchronized.
St	Stratum setting for the peer.

TABLE 7 NTP associations command output descriptions

Field	Description
when	Time, in seconds, since last NTP packet was received from peer.
poll	Polling interval (seconds).
reach	Peer reachability (bit string, in octal).
delay	Round-trip delay to peer, in milliseconds.
offset	Relative time difference between a peer clock and a local clock, in milliseconds.
disp	Dispersion.

Displaying NTP associations details

Use the **show ntp associations detail** command to display all the NTP servers and peers association information.

```
Brocade# show ntp association detail
2001:1:99:30::1 configured server, sys peer, stratum 3
ref ID 204.235.61.9, time d288dc3b.f2a17891 (10:23:55.4070668433 Pacific Tue Dec
06 2011)
our mode client, peer mode server, our poll intvl 10, peer poll intvl 10,
root delay 0.08551025 msec, root disp 0.09309387, reach 17, root dist 0.17668502
delay 0.69961487 msec, offset -13.49459670 msec, dispersion 17.31550718,
precision 2**-16, version 4
org time d288df70.a91de561 (10:37:36.2837308769 Pacific Tue Dec 06 2011)
rcv time d288df70.a0c8d19e (10:37:36.2697515422 Pacific Tue Dec 06 2011)
xmt time d288df70.a086e4de (10:37:36.2693194974 Pacific Tue Dec 06 2011)
filter delay 1.7736 0.9933 0.8873 0.6699 0.7709 0.7712 0.7734 6.7741
filter offset -17.9936 33.0014 -13.6604 -13.4494 -14.4481 -16.4453 -18.4423
-22.0025
filter disp 15.6660 0.0030 17.7730 17.7700 17.6670 17.6640 17.6610 16.6635
filter epoch 55824 56866 55686 55688 55690 55692 55694 55759
```

Use the **show ntp associations detail ipv4-address | ipv6-address** command to display the NTP servers and peers association information for a specific IP address.

```
Brocade# show ntp association detail 1.99.40.1
1.99.40.1 configured server, candidate, stratum 3
ref ID 216.45.57.38, time d288de7d.690ca5c7 (10:33:33.1762436551 Pacific Tue Dec
06 2011)
our mode client, peer mode server, our poll intvl 10, peer poll intvl 10,
root delay 0.02618408 msec, root disp 0.10108947, reach 3, root dist 0.23610585
delay 0.92163588 msec, offset 60.77749188 msec, dispersion 70.33842156,
precision 2**-16, version 4
org time d288defa.b260a71f (10:35:38.2992678687 Pacific Tue Dec 06 2011)
rcv time d288defa.a2efbd41 (10:35:38.2733620545 Pacific Tue Dec 06 2011)
xmt time d288defa.a2ae54f8 (10:35:38.2729334008 Pacific Tue Dec 06 2011)
filter delay 0.000 6.7770 6.7773 6.7711 6.7720 6.7736 6.7700 0.9921
filter offset 0.000 19.0047 19.1145 19.2245 19.3313 17.4410 15.4463 60.7777
filter disp 16000.000 16.0005 15.9975 15.9945 15.9915 15.8885 15.8855 0.0030
filter epoch 55683 55683 55685 55687 55689 55691 55693 56748
```

Syntax: show ntp association detail *ipv4-address* | *ipv6-address*

Table 7 **show ntp associations detail** command output descriptions

TABLE 8 NTP associations detail command output descriptions

Field	Description
server	Indicates server is statically configured.
symmetric active peer	Indicates peer is statically configured.
symmetric passive peer	Indicates peer is dynamically configured.
sys_peer	This peer is the system peer
candidate	This peer is chosen as candidate in the combine algorithm.
reject	This peer is rejected by the selection algorithm
falsetick	This peer is dropped as falseticker by the selection algorithm
outlyer	This peer is dropped as outlyer by the clustering algorithm
Stratum	Stratum number
ref ID	IPv4 address or hash of IPv6 address of the upstream time server to which the peer is synchronized.
Time	Last time stamp that the peer received from its master.
our mode	This system's mode relative to peer (active/passive/client/server/bdcast/bdcast client).
peer mode	Mode of peer relative to this system.
our poll intvl	This system's poll interval to this peer.
peer poll intvl	Poll interval of peer to this system
root delay	The delay along path to root (the final stratum 1 time source).
root disp	Dispersion of path to root.
reach peer	The peer reachability (bit string in octal).
Delay	Round-trip delay to peer.
offset	Offset of a peer clock relative to this clock.
Dispersion	Dispersion of a peer clock.
precision	Precision of a peer clock.
version	Peer NTP version number.
org time	Originate time stamp of the last packet.
rcv time	Receive time stamp of the last packet.
xmt time	Transmit time stamp of the last packet.
filter delay	Round-trip delay in milliseconds of last 8 samples.
filter offset	Clock offset in milliseconds of last 8 samples.
filter error	Approximate error of last 8 samples.

Configuration Examples

The following sections list configuration examples to configure the Brocade device.

NTP server and client mode configuration

Sample CLI commands to configure the Brocade device in NTP server and client modes.

```
Brocade(config-ntp)# server 10.1.2.3 minpoll 5 maxpoll 10
Brocade(config-ntp)# server 11::1/64
Brocade(config-ntp)# peer 10.100.12.18
Brocade(config-ntp)# peer 10.100.12.20
Brocade(config-ntp)# peer 10.100.12.67
Brocade(config-ntp)# peer 10.100.12.83
```

NTP client mode configuration

Sample CLI commands to configure the Brocade device in NTP client mode.

```
Brocade(config-ntp)# server 10.1.2.3 minpoll 5 maxpoll 10
Brocade(config-ntp)# server 11::1/24
Brocade(config-ntp)# peer 10.100.12.83
Brocade(config-ntp)# disable serve
```

NTP strict authentication configuration

Sample CLI commands to configure the Brocade device in strict authentication mode.

```
Brocade(config-ntp)# authenticate
Brocade(config-ntp)# authentication-key key-id 1 md5 key123
Brocade(config-ntp)# server 10.1.2.4 key 1
```

NTP loose authentication configuration

Sample CLI commands to configure the Brocade device in loose authentication mode. This allows some of the servers or clients to use the authentication keys.

```
Brocade(config-ntp)# authentication-key key-id 1 md5 key123
Brocade(config-ntp)# server 10.1.2.4 key 1
Brocade(config-ntp)# server 10.1.2.7
```

NTP interface context for the broadcast server or client mode

Sample CLI commands to enter the NTP interface context.

```
Brocade(config)#int management 1
Brocade(config-if-mgmt-1)#ip address 10.20.99.173/24
Brocade(config-if-mgmt-1)#ntp
Brocade(config-ntp)# ntp-interface management 1
Brocade(config-ntp-mgmt-1)# broadcast destination 10.23.45.128
Brocade(config-ntp)# ntp-interface ethernet 1/3
Brocade(config-ntp-if-e1000-1/3)# broadcast destination 10.1.1.0 key 1
Brocade(config-ntp)# ntp-interface ve 100
Brocade(config-ntp-ve-100)# broadcast destination 10.2.2.0 key 23
```

NTP broadcast client configuration

Sample CLI commands to configure the NTP broadcast client.

```
Brocade(config-ntp)# ntp-interface management 1
Brocade(config-ntp-mgmt-1)# broadcast client
Brocade(config-ntp)# ntp-interface ethernet 1/5
Brocade(config-ntp-if-e1000-1/5)# broadcast client
Brocade(config-ntp)# ntp-interface ve 100
Brocade(config-ntp-ve-100)# broadcast client
```

Basic port parameter configuration

The procedures in this section describe how to configure the port parameters shown in Table 5.

All Brocade ports are pre-configured with default values that allow the device to be fully operational at initial startup without any additional configuration. However, in some cases, changes to the port parameters may be necessary to adjust to attached devices or other network requirements.

Specifying a port address

You can specify a port address for an uplink (data) port, stacking port, or a management port.

ICX 6430 and ICX 6450

Specifying a data port

The port address format is is *stack unit*/*slot*/*port*, where:

- *stack unit*—Specifies the stack unit ID. For the ICX 6430, range is from 1 to 4. For the ICX 6450, range is from 1 to 8. If the device is not part of a stack, the stack unit ID is 1.
- *slot*—Specifies the slot number. Can be 1 or 2.
- *port*—Specifies the port number in the slot. Range is from 1 to 24 (24-port models) or 1 to 48 (48-port models).

This example shows how to specify port 2 in slot 1 of a device that is not part of a stack:

```
Brocade (config) # interface ethernet 1/1/2
```

Specifying a stacking port

The port address format is is *stack unit*/*slot*/*port*, where:

- *stack unit*—Specifies the stack unit ID. For the ICX 6430, range is from 1 to 4. For the ICX 6450, range is from 1 to 8.
- *slot*—Specifies the slot number. Stacking ports are in slot 2.
- *port*—Specifies the port number in the slot. Stacking ports are 1, 2, 3, and 4.

This example shows how to specify stacking port 3 in slot 2 of unit 3 in a stack:

```
Brocade (config) # interface ethernet 3/2/3
```

Specifying a management port

The management port number is always 1. This example shows how to specify the management port:

```
Brocade (config) # interface management 1
```

ICX 6610

Specifying a data port

The port address format is is *stack unit/slot/port*, where:

- *stack unit*—Specifies the stack unit ID. Range is from 1 to 8. If the device is not part of a stack, the stack unit ID is 1.
- *slot*—Specifies the slot number. Can be 1 or 3.
- *port*—Specifies the port number in the slot. Range is from 1 to 24 (24-port models) or 1 to 48 (48-port models).

This example shows how to specify port 2 in slot 1 of a device that is not part of a stack:

```
Brocade (config) # interface ethernet 1/1/2
```

Specifying a stacking port

The port address format is is *stack unit/slot/port*, where:

- *stack unit*—Specifies the stack unit ID. Range is from 1 to 8.
- *slot*—Specifies the slot number. Stacking ports are in slot 2.
- *port*—Specifies the port number in the slot. Dedicated stacking ports are 1, 2, 6, and 7.

This example shows how to specify stacking port 2 in slot 2 of unit 3 in a stack:

```
Brocade (config) # interface ethernet 3/2/2
```

Specifying a management port

The management port number is always 1. This example shows how to specify the management port:

```
Brocade (config) # interface management 1
```

FCX

Specifying a data port

The port address format is *stack unit/slot/port*, where:

- *stack unit*—Specifies the stack unit ID. Range is from 1 to 8. If the device is not part of a stack, the stack unit ID is 1.
- *slot*—Specifies the slot number. Can be 1 or 3.
- *port*—Specifies the port number in the slot. Range is from 1 to 24 (24-port models) or 1 to 48 (48-port models).

This example shows how to specify port 2 in slot 1 of a device that is not part of a stack:

```
Brocade (config) # interface ethernet 1/1/2
```

Specifying a stacking port

The port address format is *stack unit/slot/port*, where:

- *stack unit*—Specifies the stack unit ID. Range is from 1 to 8.
- *slot*—Specifies the slot number. Default stacking ports are in slot 2 (FCX S/S-F) and slot3 (FCX E/I).
- *port*—Specifies the port number in the slot. Default stacking ports in slot 2 and slot 3 are ports 1 and 2.

This example shows how to specify port 2 in slot 2 of unit 3 in a stack:

```
Brocade (config) # interface ethernet 3/2/2
```

Specifying a management port

The management port number is always 1. This example shows how to specify the management port:

```
Brocade (config) # interface management 1
```

FSX

Specifying a data port

The port address format is *slot/port*, where:

- *slot*—Specifies the interface slot number. Range is from 1 to 8 (FSX 800) or 1 to 16 (FSX 1600).
- *port*—Specifies the port number in the slot. Range is from 1 to 48 depending on the interface module.

This example shows how to specify port 2 in slot 1:

```
Brocade (config) # interface ethernet 1/2
```

Specifying a management port

The management port number is always 1. This example shows how to specify the management port:

```
Brocade (config) # interface management 1
```

NOTE
Stacking is not supported on FSX devices.

Assigning port names

You can assign text strings as port names, which help you identify ports with meaningful names. You can assign port names to individual ports or to a group of ports. You can assign a port name to physical ports, virtual interfaces, and loopback interfaces.

Assigning a port name

To assign a name to a port, enter commands such as the following:

```
Brocade(config)# interface ethernet 2
Brocade(config-if-e1000-2)# port-name Marsha
```

Syntax: **port-name** *text*

The *text* parameter is an alphanumeric string. The name can be up to 255 characters long. The name can contain blanks. You do not need to use quotation marks around the string, even when it contains blanks. The port name can contain special characers as well, but the percentage character (%), if it appears at the end of the port name, is dropped.

Assigning the same name to multiple ports

To assign a name to a range of ports, enter commands such as the following:

```
Brocade (config)# interface ethernet 1/1/1 to 1/1/10
Brocade (config-mif-1/1/1-1/1/10)# port-name connected-to-the nearest device
```

Syntax: [**no**] **Port-name** *text*

To remove the assigned port name, use **no** form of the command.

The *text* parameter is an alphanumeric string, up to 255 characters long. The name can contain blanks. You do not need to use quotation marks around the string, even when it contains blanks.

You can also specify the individual ports, separated by space.

To assign a name to multiple specific ports, enter commands such as the following:

```
Brocade (config)# interface ethernet 1/1/1 ethernet 1/1/5 ethernet 1/1/7
Brocade (config-mif-1/1/1, 1/1/5, 1/1/7)# port-name connected-to-the nearest
device
```

Displaying the port name for an interface

You can use the **show interfaces brief** command to display the name assigned to the port. If any of the ports have long port names, they are truncated. To show full port names, use the show interfaces brief wide command.

```
Brocade# show interfaces brief

Port    Link    State   Dupl Speed Trunk Tag Pvid Pri MAC            Name
1/1/23  Up      Forward Full 1G    None  No  1    0   748e.f82d.7a16 connected-
1/1/47  Up      Forward Full 1G    None  No  1    0   748e.f82d.7a2e
mgmt1   Up      None    Full 1G    None  No  None 0   748e.f82d.7a00
```

In this output, the port name for inteface 1/1/23 is truncated.

Use the **show interface brief wide** command to avoid truncating long port names.

To display the complete port name for an interface, enter the following command.

```
Brocade# show interface brief wide

Port    Link    State   Dupl Speed Trunk Tag Pvid Pri MAC            Name
1/1/23  Up      Forward Full 1G    None  No  1    0   748e.f82d.7a16 connected-
to-the nearest device
1/1/47  Up      Forward Full 1G    None  No  1    0   748e.f82d.7a2e
mgmt1   Up      None    Full 1G    None  No  None 0   748e.f82d.7a00
```

Syntax: **show interface brief wide** [**ethernet** *stack-unit/slot/port* | **loopback** | **management** | **slot** | **tunnel** | **ve**]

The **ethernet** *stack-unit/slot/port* parameter specifies the Ethernet port for which you want to display the interface information.

The **loopback** option specifies the loopback port for which you want to display the interface information.

The **management** option specifies the management port for which you want to display the interface information.

The **slot** option specifies all the ports in a slot for which you want to display the interface information.

The **tunnel** option specifies the tunnel port for which you want to display the interface information.

The **ve** option specifies the virtual routing (VE) port for which you want to display the interface information.

Table 9 describes the output parameters of the **show interface brief wide** command.

TABLE 9 Output parameters of the show interface brief wide command

Field	Description
Port	Specifies the port number.
Link	Specifies the link state.
Port-State	Specifies the current port state.
Speed	Specifies the link speed.
Tag	Specifies if the port is tagged or not.
Pvid	Specifies the port VLAN ID.
Pri	Specifies the priority.
MAC	Specifies the MAC address.
Name	Specifies the port name.

To display the complete port name for an Ethernet interface, enter a command such as the following.

```
Brocade# show interface brief wide ethernet 1/1/23

PPort    Link    State   Dupl Speed Trunk Tag Pvid Pri MAC            Name
1/1/23   Up      Forward Full 1G    None  No  1    0   748e.f82d.7a16  connected-
to-FCX
```

Syntax: **show interface brief wide ethernet** *stack-unit/slot/port*

For more information about field descriptions of the command output, refer Table 9.

Port speed and duplex mode modification

The Gigabit Ethernet copper ports are designed to auto-sense and auto-negotiate the speed and duplex mode of the connected device. If the attached device does not support this operation, you can manually enter the port speed to operate at either 10, 100, or 1000 Mbps. The default and recommended setting is 10/100/1000 auto-sense.

NOTE
You can modify the port speed of copper ports only; this feature does not apply to fiber ports.

NOTE
For optimal link operation, copper ports on devices that do not support 803.3u must be configured with like parameters, such as speed (10,100,1000), duplex (half, full), MDI/MDIX, and Flow Control.

Port speed and duplex mode configuration syntax

The following commands change the port speed of copper interface 8 on a FastIron from the default of 10/100/1000 auto-sense, to 100 Mbps operating in full-duplex mode.

```
Brocade(config)# interface ethernet 8
Brocade(config-if-e1000-8)# speed-duplex 100-full
```

Syntax: **speed-duplex** *value*

where *value* can be one of the following:

- 10-full – 10 Mbps, full duplex
- 10-half – 10 Mbps, half duplex
- 100-full – 100 Mbps, full duplex
- 100-half – 100 Mbps, half duplex
- 1000-full-master – 1 Gbps, full duplex master
- 1000-full-slave – 1 Gbps, full duplex slave
- auto – auto-negotiation

The default is **auto** (auto-negotiation).

Use the **no** form of the command to restore the default.

NOTE
On FastIron devices, when setting the speed and duplex-mode of an interface to 1000-full, configure one side of the link as master (**1000-full-master**) and the other side as slave (**1000-full-slave**).

NOTE
On Brocade ICX 6610 devices, after you remove 10 Gbps speed from the running configuration, plugging in a 1G optic SFP transceiver into a 10 Gbps port causes the software to fail to revert the ports back from the default 10G LRM mode to 1 Gbps speed. Remove the 1G SFP transceiver and plug in the 10G optic SFP+transceiver so that the Brocade ICX 6610 devices go into default 10 Gbps LRM mode.

Enabling auto-negotiation maximum port speed advertisement

NOTE
For optimal link operation, link ports on devices that do not support 802.3u must be configured with like parameters, such as speed (10,100,1000), duplex (half, full), MDI/MDIX, and Flow Control.

Maximum Port speed advertisement is an enhancement to the auto-negotiation feature, a mechanism for accommodating multi-speed network devices by automatically configuring the highest performance mode of inter-operation between two connected devices.

Maximum port speed advertisement enables you to configure an auto-negotiation maximum speed that Gbps copper ports on the Brocade device will advertise to the connected device. You can configure a port to advertise a maximum speed of either 100 Mbps or 10 Mbps. When the maximum port speed advertisement feature is configured on a port that is operating at 100 Mbps maximum speed, the port will advertise 10/100 Mbps capability to the connected device. Similarly, if a port is configured at 10 Mbps maximum speed, the port will advertise 10 Mbps capability to the connected device.

The maximum port speed advertisement feature operates dynamically at the physical link layer between two connected network devices. They examine the cabling conditions and the physical capabilities of the remote link, then configure the speed of the link segment according to the highest physical-layer technology that both devices can accommodate.

The maximum port speed advertisement feature operates dynamically at the physical link layer, independent of logical trunk group configurations. Although Brocade recommends that you use the same cable types and auto-negotiation configuration on all members of a trunk group, you could utilize the auto-negotiation features conducive to your cabling environment. For example, in certain circumstances, you could configure each port in a trunk group to have its own auto-negotiation maximum port speed advertisement.

Maximum port speed application notes

- The maximum port speed advertisement works only when auto-negotiation is enabled (CLI command **speed-duplex auto**). If auto-negotiation is OFF, the device will reject the maximum port speed advertisement configuration.
- When the maximum port speed advertisement is enabled on a port, the device will reject any configuration attempts to set the port to a forced speed mode (100 Mbps or 1000 Mbps).

Configuring maximum port speed advertisement

NOTE
This is not supported in ICX devices.

To configure a maximum port speed advertisement of 10 Mbps on a port that has auto-negotiation enabled, enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)# link-config gig copper autoneg-control 10m ethernet 1
```

To configure a maximum port speed advertisement of 100 Mbps on a port that has auto-negotiation enabled, enter the following command at the Global CONFIG level of the CLI.

```
Brocade(config)# link-config gig copper autoneg-control 100m ethernet 2
```

Syntax: [**no**] **link-config gig copper autoneg-control 10m** | **100m ethernet** *port* [**ethernet** [*port*]

You can enable maximum port speed advertisement on one or two ports at a time.

To disable maximum port speed advertisement after it has been enabled, enter the **no** form of the command.

Modifying port duplex mode

You can manually configure a 10/100 Mbps port to accept either full-duplex (bi-directional) or half-duplex (uni-directional) traffic.

NOTE
You can modify the port duplex mode of copper ports only. This feature does not apply to fiber ports.

Port duplex mode and port speed are modified by the same command.

Port duplex mode configuration syntax

To change the port speed of interface 8 from the default of 10/100/1000 auto-sense to 10 Mbps operating at full-duplex, enter the following.

```
Brocade(config)# interface ethernet 8
Brocade(config-if-e1000-8)# speed-duplex 10-full
```

Syntax: **speed-duplex** *value*

The *value* can be one of the following:

- 10-full
- 10-half
- 100-full
- 100-half
- auto (default)

MDI and MDIX configuration

Brocade devices support automatic Media Dependent Interface (MDI) and Media Dependent Interface Crossover (MDIX) detection on all Gbps Ethernet Copper ports.

MDI/MDIX is a type of Ethernet port connection using twisted pair cabling. The standard wiring for end stations is MDI, whereas the standard wiring for hubs and switches is MDIX. MDI ports connect to MDIX ports using straight-through twisted pair cabling. For example, an end station connected to a hub or a switch uses a straight-through cable. MDI-to-MDI and MDIX-to-MDIX connections use crossover twisted pair cabling. So, two end stations connected to each other, or two hubs or switches connected to each other, use crossover cable.

The auto MDI/MDIX detection feature can automatically correct errors in cable selection, making the distinction between a straight-through cable and a crossover cable insignificant.

MDI and MDIX configuration notes

- This feature applies to copper ports only.
- The **mdi-mdix mdi** and **mdi-mdix mdix** commands work independently of auto-negotiation. Thus, these commands work whether auto-negotiation is turned ON or OFF.

MDI and MDIX configuration syntax

The auto MDI/MDIX detection feature is enabled on all Gbps copper ports by default. For each port, you can disable auto MDI/MDIX, designate the port as an MDI port, or designate the port as an MDIX port.

To turn off automatic MDI/MDIX detection and define a port as an MDI only port.

```
Brocade(config-if-e1000-2)# mdi-mdix mdi
```

To turn off automatic MDI/MDIX detection and define a port as an MDIX only port.

```
Brocade(config-if-e1000-2)# mdi-mdix mdix
```

To turn on automatic MDI/MDIX detection on a port that was previously set as an MDI or MDIX port.

```
Brocade(config-if-e1000-2)# mdi-mdix auto
```

Syntax: **mdi-mdix** *mdi* | *mdix* | *auto*

After you enter the **mdi-mdix** command, the Brocade device resets the port and applies the change.

To display the MDI/MDIX settings, including the configured value and the actual resolved setting (for mdi-mdix auto), enter the command **show interface** at any level of the CLI.

Disabling or re-enabling a port

A port can be made inactive (disable) or active (enable) by selecting the appropriate status option. The default value for a port is enabled.

To disable port 8 of a Brocade device, enter the following.

```
Brocade(config)# interface ethernet 8
Brocade(config-if-e1000-8)# disable
```

You also can disable or re-enable a virtual interface. To do so, enter commands such as the following.

```
Brocade(config)# interface ve v1
Brocade(config-vif-1)# disable
```

Syntax: **disable**

To re-enable a virtual interface, enter the **enable** command at the Interface configuration level. For example, to re-enable virtual interface v1, enter the **enable** command.

```
Brocade(config-vif-1)# enable
```

Syntax: **enable**

Flow control configuration

Flow control (802.3x) is a QoS mechanism created to manage the flow of data between two full-duplex Ethernet devices. Specifically, a device that is oversubscribed (is receiving more traffic than it can handle) sends an 802.3x PAUSE frame to its link partner to temporarily reduce the amount of data the link partner is transmitting. Without flow control, buffers would overflow, packets would be dropped, and data retransmission would be required.

All FastIron devices support *asymmetric* flow control, meaning they can receive PAUSE frames but cannot transmit them. In addition, FCX and ICX devices also support *symmetric* flow control, meaning they can both receive and transmit 802.3x PAUSE frames. For details about symmetric flow control, refer to "Symmetric flow control on FCX and ICX devices" on page 50.

Flow control configuration notes

- Auto-negotiation of flow control is not supported on 10 Gbps ports, fiber ports, and copper or fiber combination ports.
- When any of the flow control commands are applied to a port that is up, the port will be disabled and re-enabled.
- For 10 Gbps ports, the **show interface** *port* display shows Flow Control is enabled or Flow Control is disabled, depending on the configuration.
- When flow-control is enabled, the hardware can only advertise PAUSE frames. It does not advertise Asym.

Disabling or re-enabling flow control

You can configure the Brocade device to operate with or without flow control. Flow control is enabled by default globally and on all full-duplex ports. You can disable and re-enable flow control at the Global CONFIG level for all ports. When enabled globally, you can disable and re-enable flow control on individual ports.

To disable flow control, enter the **no flow-control** command.

```
Brocade(config)# no flow-control
```

To turn the feature back on, enter the **flow-control** command.

```
Brocade(config)# flow-control
```

Syntax: [**no**] **flow-control**

NOTE
For optimal link operation, link ports on devices that do not support 803.3u must be configured with like parameters, such as speed (10,100,1000), duplex (half, full), MDI/MDIX, and Flow Control.

Negotiation and advertisement of flow control

By default, when flow control is enabled globally and auto-negotiation is ON, flow control is enabled and advertised on 10/100/1000M ports. If auto-negotiation is OFF or if the port speed was configured manually, then flow control is not negotiated with or advertised to the peer. For details about auto-negotiation, refer to "Port speed and duplex mode modification" on page 43.

To disable flow control capability on a port, enter the following commands.

```
Brocade(config)# interface ethernet 0/1/21
Brocade(config-if-e1000-0/1/21)# no flow-control
```

To enable flow control negotiation, enter the following commands.

```
Brocade(config)# interface ethernet 0/1/21
Brocade(config-if-e1000-0/1/21)# flow-control neg-on
```

Syntax: [**no**] **flow-control** [**neg-on**]

- **flow-control** [default] - Enable flow control, flow control negotiation, and advertise flow control
- **no flow-control neg-on** - Disable flow control negotiation
- **no flow-control** - Disable flow control, flow control negotiation, and advertising of flow control

After flow control negotiation is enabled using the **flow-control neg-on** command option, flow control is enabled or disabled depending on the peer advertisement.

Commands may be entered in IF (single port) or MIF (multiple ports at once) mode.

Example

```
Brocade(config)# interface ethernet 0/1/21
Brocade(config-if-e1000-0/1/21)# no flow-control
```

This command disables flow control on port 0/1/21.

```
Brocade(config)# interface ethernet 0/1/11 to 0/1/15
Brocade(config-mif-0/1/11-0/1/15)# no flow-control
```

This command disables flow control on ports 0/1/11 to 0/1/15.

Displaying flow-control status

The **show interface** *port* command displays configuration, operation, and negotiation status where applicable.

For example, on a FastIron Stackable device, issuing the command for 10/100/1000M port 0/1/21 displays the following output.

```
Brocade# show interfaces ethernet 0/1/21
GigabitEthernet0/1/21 is up, line protocol is up
  Port up for 30 minutes 20 seconds
  Hardware is GigabitEthernet, address is 0000.0004.4014 (bia 0000.0004.4014)
  Configured speed auto, actual 100Mbit, configured duplex fdx, actual fdx
  Configured mdi mode AUTO, actual MDIX
  Member of L2 VLAN ID 1, port is untagged, port state is LISTENING
  BPDU Guard is disabled, Root Protect is disabled
  STP configured to ON, priority is level0
  Flow Control is config enabled, oper enabled, negotiation disabled
  Mirror disabled, Monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  Inter-Packet Gap (IPG) is 96 bit times
  300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  300 second output rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts, 0 multicasts, 0 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 runts, 0 giants
  5 packets output, 320 bytes, 0 underruns
  Transmitted 0 broadcasts, 5 multicasts, 0 unicasts
  0 output errors, 0 collisions
```

NOTE
The port up/down time is required only for physical ports and not for loopback/ ve/ tunnel ports.

Issuing the **show interface** *port* command on a FSX device displays the following output:

```
Brocade# show interface ethernet 18/1
GigabitEthernet18/1 is up, line protocol is up
  Port up for 50 seconds
  Hardware is GigabitEthernet, address is 0000.0028.0600 (bia 0000.0028.0798)
  Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
  Configured mdi mode AUTO, actual MDIX
  Member of 4 L2 VLANs, port is tagged, port state is FORWARDING
  BPDU guard is Disabled, ROOT protect is Disabled
  Link Error Dampening is Disabled
  STP configured to ON, priority is level0, flow control enabled
  Flow Control is config enabled, oper enabled, negotiation disabled
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  IPG MII 96 bits-time, IPG GMII 96 bits-time
  IP MTU 1500 bytes, encapsulation ethernet
  300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  300 second output rate: 848 bits/sec, 0 packets/sec, 0.00% utilization
  0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts, 0 multicasts, 0 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 runts, 0 giants
  10251 packets output, 1526444 bytes, 0 underruns
  Transmitted 1929 broadcasts, 8293 multicasts, 29 unicasts
  0 output errors, 0 collisions
```

The line highlighted in bold will resemble one of the following, depending on the configuration:

- If flow control negotiation is enabled (and a neighbor advertises "Pause-Not Capable"), the display shows:

  `Flow Control is config enabled, oper disabled, negotiation enabled`

- If flow control negotiation is enabled (and a neighbor advertises "Pause-Capable"), the display shows:

  `Flow Control is config enabled, oper enabled, negotiation enabled`

- If flow control is enabled, and flow control negotiation is disabled, the display shows:

  `Flow Control is config enabled, oper enabled, negotiation disabled`

- If flow control is disabled, the display shows:

  `Flow control is config disabled, oper disabled`

Symmetric flow control on FCX and ICX devices

In addition to *asymmetric* flow control, FCX and ICX devices support *symmetric* flow control, meaning they can both receive and transmit 802.3x PAUSE frames.

By default on FCX devices, packets are dropped from the end of the queue at the egress port (tail drop mode), when the maximum queue limit is reached. Conversely, when symmetric flow control is enabled, packets are guaranteed delivery since they are managed at the ingress port and no packets are dropped.

Symmetric flow control addresses the requirements of a lossless service class in an Internet Small Computer System Interface (iSCSI) environment. It is supported on FCX and ICX standalone units as well as on all FCX and ICX units in a traditional stack.

About XON and XOFF thresholds

An 802.3x PAUSE frame is generated when the buffer limit at the ingress port reaches or exceeds the port's upper watermark threshold (XOFF limit). The PAUSE frame requests that the sender stop transmitting traffic for a period of time. The time allotted enables the egress and ingress queues to be cleared. When the ingress queue falls below the port's lower watermark threshold (XON limit), an 802.3x PAUSE frame with a quanta of 0 (zero) is generated. The PAUSE frame requests that the sender resume sending traffic normally.

Each 1G and 10G port is configured with a default total number of buffers as well as a default XOFF and XON threshold. The defaults are different for 1G ports versus 10G ports. Also, the default XOFF and XON thresholds are different for jumbo mode versus non-jumbo mode. The defaults are shown in Table 10.

TABLE 10 XON and XOFF default thresholds

	Limit when Jumbo disabled / % of buffer limit	Limit when Jumbo enabled / % of buffer limit
1G ports		
Total buffers	272	272
XOFF	240 / 91%	216 / 82%
XON	200 / 75%	184 / 70%
10G ports		
Total buffers	416	416
XOFF	376 / 91%	336 / 82%
XON	312 / 75%	288 / 70%

If necessary, you can change the total buffer limits and the XON and XOFF default thresholds. Refer to "Changing the total buffer limits" on page 53 and "Changing the XON and XOFF thresholds" on page 52, respectively.

Configuration notes and feature limitations for symmetric flow control

Note the following configuration notes and feature limitations before enabling symmetric flow control.

- Symmetric flow control is supported on FCX and ICX devices only. It is not supported on other FastIron models.
- Symmetric flow control is supported on all 1G and 10G data ports on FCX and ICX devices.
- Symmetric flow control is not supported on stacking ports or across units in a stack.
- To use this feature, 802.3x flow control must be enabled globally and per interface on FCX and ICX devices. By default, 802.3x flow control is enabled, but can be disabled with the **no flow-control** command.
- The following QoS features are not supported together with symmetric flow control:
  - Dynamic buffer allocation (CLI commands **qd-descriptor** and **qd-buffer**)
  - Buffer profiles (CLI command **buffer-profile port-region**)
  - DSCP-based QoS (CLI command **trust dscp**)

NOTE
Although the above QoS features are not supported with symmetric flow control, the CLI will still accept these commands. The last command issued will be the one placed into effect on the device. For example, if **trust dscp** is enabled after **symmetric-flow-control** is enabled, symmetric flow control will be disabled and trust dscp will be placed into effect. Make sure you do not enable incompatible QoS features when symmetric flow control is enabled on the device.

- Head of Line (HOL) blocking may occur when symmetric flow control is enabled. This means that a peer can stop transmitting traffic streams unrelated to the congestion stream.

Enabling and disabling symmetric flow control

By default, symmetric flow control is disabled and tail drop mode is enabled. However, because flow control is enabled by default on all full-duplex ports, these ports will always honor received 802.3x Pause frames, whether or not symmetric flow control is enabled.

To enable symmetric flow control globally on all full-duplex data ports of a standalone unit, enter the **symmetric-flow-control enable** command.

```
Brocade(config)# symmetric-flow-control enable
```

To enable symmetric flow control globally on all full-duplex data ports of a particular unit in a traditional stack, enter the **symmetric-flow-control enable** *stack-unit* command.

```
Brocade(config)# symmetric-flow-control enable unit 4
```

Syntax: [**no**] **symmetric-flow-control enable** [**unit** *stack-unit*]

The *stack-unit* parameter specifies one of the units in a stacking system.
Master/Standby/Members are examples of a stack-unit

To disable symmetric flow control once it has been enabled, use the **no** form of the command.

Changing the XON and XOFF thresholds

This section describes how to change the XON and XOFF thresholds described in "About XON and XOFF thresholds" on page 51.

To change the thresholds for all 1G ports, enter a command such as the following.

```
Brocade(config)# symmetric-flow-control set 1 xoff 91 xon 75
```

To change the thresholds for all 10G ports, enter a command such as the following.

```
Brocade(config)# symmetric-flow-control set 2 xoff 91 xon 75
```

In the above configuration examples, when the XOFF limit of 91% is reached or exceeded, the Brocade device will send PAUSE frames to the sender telling it to stop transmitting data temporarily. When the XON limit of 75% is reached, the Brocade device will send PAUSE frames to the sender telling it to resume sending data.

Syntax: **symmetric-flow-control set 1** | **2 xoff** % **xon** %

symmetric-flow-control set 1 sets the XOFF and XON limits for 1G ports.

symmetric-flow-control set 2 sets the XOFF and XON limits for 10G ports.

For **xoff** %, the % minimum value is 60% and the maximum value is 95%.

For **xon** %, the % minimum value is 50% and the maximum value is 90%.

Use the **show symmetric** command to view the default or configured XON and XOFF thresholds. Refer to "Displaying symmetric flow control status" on page 53.

Changing the total buffer limits

This section describes how to change the total buffer limits described in "About XON and XOFF thresholds" on page 51. You can change the limits for all 1G ports and for all 10G ports.

To change the total buffer limit for all 1G ports, enter a command such as the following.

```
Brocade(config)# symmetric-flow-control set 1 buffers 320
Total buffers modified, 1G: 320, 10G: 128
```

To change the total buffer limit for all 10G ports, enter a command such as the following.

```
Brocade(config)# symmetric-flow-control set 2 buffers 128
Total buffers modified, 1G: 320, 10G: 128
```

Syntax: **symmetric-flow-control set 1 | 2 buffers** *value*

symmetric-flow-control set 1 buffers *value* sets the total buffer limits for 1G ports. The default *value* is 272. You can specify a number from 64 – 320.

symmetric-flow-control set 2 buffers *value* sets the total buffer limits for 10G ports. The default *value* is 416. You can specify a number from 64 – 1632.

Use the **show symmetric** command to view the default or configured total buffer limits. Refer to "Displaying symmetric flow control status" on page 53.

Displaying symmetric flow control status

The **show symmetric-flow-control** command displays the status of symmetric flow control as well as the default or configured total buffer limits and XON and XOFF thresholds.

```
Brocade(config)# show symmetric
Symmetric Flow Control Information:
-----------------------------------
Symmetric Flow Control is enabled on units: 2 3
Buffer parameters:
1G Ports:
        Total Buffers : 272
        XOFF Limit    : 240(91%)
        XON Limit     : 200(75%)
10G Ports:
        Total Buffers : 416
        XOFF Limit    : 376(91%)
        XON Limit     : 312(75%)
```

Syntax: **show symmetric-flow-control**

PHY FIFO Rx and Tx depth configuration

PHY devices on Brocade devices contain transmit and receive synchronizing FIFOs to adjust for frequency differences between clocks. The **phy-fifo-depth** command allows you to configure the depth of the transmit and receive FIFOs. There are 4 settings (0-3) with 0 as the default. A higher setting indicates a deeper FIFO.

The default setting works for most connections. However, if the clock differences are greater than the default will handle, CRCs and errors will begin to appear on the ports. Raising the FIFO depth setting will adjust for clock differences.

Brocade recommends that you disable the port before applying this command, and re-enable the port. Applying the command while traffic is flowing through the port can cause CRC and other errors for any packets that are actually passing through the PHY while the command is being applied.

Syntax: [**no**] **phy-fifo-depth** *setting*

- *setting* is a value between 0 and 3. (0 is the default.)

This command can be issued for a single port from the IF config mode or for multiple ports from the MIF config mode.

NOTE
Higher settings give better tolerance for clock differences with the partner phy, but may marginally increase latency as well.

Interpacket Gap (IPG) on a FastIron X Series switch

IPG is the time delay, in bit time, between frames transmitted by the device. You configure IPG at the interface level. The command you use depends on the interface type on which IPG is being configured.

The default interpacket gap is 96 bits-time, which is 9.6 microseconds for 10 Mbps Ethernet, 960 nanoseconds for 100 Mbps Ethernet, 96 nanoseconds for 1 Gbps Ethernet, and 9.6 nanoseconds for 10 Gbps Ethernet.

IPG on a FastIron X series switch configuration notes

- The CLI syntax for IPG differs on FastIron X Series devices compared to FastIron Stackable devices. This section describes the configuration procedures for FastIron X Series devices. For FastIron Stackable devices, refer to "IPG on FastIron Stackable devices" on page 55.
- IPG configuration commands are based on "port regions". All ports within the same port region should have the same IPG configuration. If a port region contains two or more ports, changes to the IPG configuration for one port are applied to all ports in the same port region. When you enter a value for IPG, the CLI displays the ports to which the IPG configuration is applied.

  Example

  ```
  Brocade(config-if-e1000-7/1)# ipg-gmii 120
  IPG 120(112) has been successfully configured for ports 7/1 to 7/12
  ```

- When you enter a value for IPG, the device applies the closest valid IPG value for the port mode to the interface. For example, if you specify 120 for a 1 Gbps Ethernet port in 1 Gbps mode, the device assigns 112 as the closest valid IPG value to program into hardware.

Configuring IPG on a Gbps Ethernet port

On a Gbps Ethernet port, you can configure IPG for 10/100 mode and for Gbps Ethernet mode.

10/100M mode

To configure IPG on a Gbps Ethernet port for 10/100M mode, enter the following command.

```
Brocade(config)# interface ethernet 7/1
Brocade(config-if-e1000-7/1)# ipg-mii 120
IPG 120(120) has been successfully configured for ports 7/1 to 7/12
```

Syntax: [**no**] **ipg-mii** *bit time*

Enter 12-124 for *bit time*. The default is 96 bit time.

1G mode

To configure IPG on a Gbps Ethernet port for 1-Gbps Ethernet mode, enter commands such as the following.

```
Brocade(config)# interface ethernet 7/1
Brocade(config-if-e1000-7/1)# ipg-gmii 120
IPG 120(112) has been successfully configured for ports 0/7/1 to 7/12
```

Syntax: [**no**] **ipg-gmii** *bit time*

Enter 48 - 112 for *bit time*. The default is 96 bit time.

Configuring IPG on a 10 Gbps Ethernet interface

To configure IPG on a 10 Gbps Ethernet interface, enter commands such as the following.

```
Brocade(config)# interface ethernet 9/1
Brocade(config-if-e10000-9/1)# ipg-xgmii 120
IPG 120(128) has been successfully configured for port 9/1
```

Syntax: [**no**] **ipg-xgmii** *bit time*

Enter 96-192 for *bit time*. The default is 96 bit time.

IPG on FastIron Stackable devices

On FCX and ICX devices, you can configure an IPG for each port. An IPG is a configurable time delay between successive data packets.

You can configure an IPG with a range from 48-120 bit times in multiples of 8, with a default of 96. The IPG may be set from either the interface configuration level or the multiple interface level.

IPG configuration notes

- The CLI syntax for IPG differs on FastIron Stackable devices compared to FastIron X Series devices. This section describes the configuration procedures for FastIron Stackable devices. For FastIron X Series devices, refer to "Interpacket Gap (IPG) on a FastIron X Series switch" on page 54.
- When an IPG is applied to a trunk group, it applies to all ports in the trunk group. When you are creating a new trunk group, the IPG setting on the primary port is automatically applied to the secondary ports.
- This feature is supported on 10/100/1000M ports.

Configuring IPG on a 10/100/1000M port

To configure an IPG of 112 on Ethernet interface 0/1/21, for example, enter the following command.

```
Brocade(config)# interface ethernet 0/1/21
Brocade(config-if-e1000-0/1/21)# ipg 112
```

For multiple interface levels, to configure IPG for ports 0/1/11 and 0/1/14 through 0/1/17, enter the following commands.

```
Brocade(config)# interface ethernet 0/1/11 ethernet 0/1/14 to 0/1/17
Brocade(config-mif-0/1/11,0/1/14-0/1/17)# ipg 104
```

Syntax: [no] ipg *value*

For *value*, enter a number in the range from 48-120 bit times in multiples of 8. The default is 96.

As a result of the above configuration, the output from the show interface Ethernet 0/1/21 command is as follows.

```
Brocade# show interfaces ethernet 0/1/21
GigabitEthernet 0/1/21 is up, line protocol is up
  Port up for 40 seconds
  Hardware is GigabitEthernet, address is 0000.0004.4014 (bia 0000.0004.4014)
  Configured speed auto, actual 100Mbit, configured duplex fdx, actual fdx
  Configured mdi mode AUTO, actual MDIX
  Member of L2 VLAN ID 1, port is untagged, port state is FORWARDING
  BPDU Guard is disabled, Root Protect is disabled
  STP configured to ON, priority is level0
  Flow Control is config enabled, oper enabled, negotiation disabled
  Mirror disabled, Monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  Inter-Packet Gap (IPG) is 112 bit times
  IP MTU 10222 bytes
  300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  300 second output rate: 248 bits/sec, 0 packets/sec, 0.00% utilization
  0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts, 0 multicasts, 0 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 runts, 0 giants
  80 packets output, 5120 bytes, 0 underruns
  Transmitted 0 broadcasts, 80 multicasts, 0 unicasts
  0 output errors, 0 collisions
```

Enabling and disabling support for 100BaseTX

For FastIron X Series devices, you can configure a 1000Base-TX SFP (part number E1MG-TX) to operate at a speed of 100 Mbps. To do so, enter the **100-tx** command at the Interface level of the CLI.

```
Brocade(config-if-e1000-11)# 100-tx
```

After the link is up, it will be in 100M/full-duplex mode, as shown in the following example.

```
Brocade# show interface brief ethernet 11
Port Link  State     Dupl    Speed Trunk  Tag   Priori     MAC Name
11    Up   Forward   Full    100M  None   No   level10    0000.0013.c74b
```

The **show media** command will display the SFP transceiver as **1G M-TX**.

Syntax: [no] 100-tx

To disable support, enter the **no** form of the command.

100BaseTX configuration notes

- This feature requires that autonegotiation be enabled on the other end of the link.
- Although combo ports (ports 1 – 4) on Hybrid Fiber (HF) models support the 1000Base-TX SFP, they cannot be configured to operate at 100 Mbps. The 100 Mbps operating speed is supported only with non-combo ports (ports 5-24).
- The FCX624S-F is the only FCX model that supports the 1000Base-TX SFP module, and only on the non-combo ports (ports 5-24). The FCX624S-F does not have a specific command to enable the 1000Base-TX SFP optic at 100 Mbps. You must manually configure it with the **speed-duplex 100-full** command. Refer to “Port speed and duplex mode configuration syntax” on page 44.
- 1000Base-TX modules must be configured individually, one interface at a time.
- 1000Base-TX modules do not support Digital Optical Monitoring.
- This module requires a Cat5 cable and uses an RJ45 connector.
- Hotswap is supported for this module when it is configured in 100M mode.

Enabling and disabling support for 100BaseFX

Some Brocade devices support 100BaseFX fiber transceivers. After you physically install a 100BaseFX transceiver, you must enter a CLI command to enable it. For information about supported SFP and SFP+ transceivers on ICX devices, refer to the following Brocade website:

http://www.brocade.com/downloads/documents/data_sheets/product_data_sheets/Optics_DS.pdf

Enabling and disabling 100BaseFX on Chassis-based and stackable devices

---

NOTE
The following procedure applies to Stackable devices and to Chassis-based 100/1000 Fiber interface modules only. The CLI syntax for enabling and disabling 100BaseFX support on these devices differs than on a Compact device. Make sure you refer to the appropriate procedures. These are not supported on ICX 6430 and ICX 6450 devices.

FastIron devices support the following types of SFPs for 100BaseFX:

- **Multimode SFP** – maximum distance is 2 kilometers
- **Long Reach (LR)** – maximum distance is 40 kilometers
- **Intermediate Reach (IR)** – maximum distance is 15 kilometers

For information about supported SFP and SFP+ transceivers on FastIron devices, refer to the following Brocade website:

> http://www.brocade.com/downloads/documents/data_sheets/product_data_sheets/Optics_DS.pdf

---

NOTE
Connect the 100BaseFX fiber transceiver *after* configuring both sides of the link. Otherwise, the link could become unstable, fluctuating between up and down states.

To enable support for 100BaseFX on an FSX fiber port or on a Stackable switch, enter commands such as the following.

```
Brocade(config)# interface ethernet 1/6
Brocade(config-if-1/6)# 100-fx
```

The above commands enable 100BaseFX on port 6 in slot 1.

Syntax: [**no**] **100-fx**

To disable 100BaseFX support on a fiber port, enter the **no** form of the command. Note that you must disable 100BaseFX support before inserting a different type of module In the same port. Otherwise, the device will not recognize traffic traversing the port.

Changing the Gbps fiber negotiation mode

The globally configured Gbps negotiation mode is the default mode for all Gbps fiber ports. You can override the globally configured default and set individual ports to the following:

NOTE
Gbps negotiation is not supported on ICX 6430 and ICX 6450 devices.

- **Negotiate-full-auto** – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information). This is the default.
- **Auto-Gbps** – The port tries to perform a handshake with the other port to exchange capability information.

- **Negotiation-off** – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.

To change the mode for individual ports, enter commands such as the following.

```
Brocade(config)# interface ethernet 1 to 4
Brocade(config-mif-1-4)# gig-default auto-gig
```

This command overrides the global setting and sets the negotiation mode to auto-Gbps for ports 1 – 4.

Syntax: **gig-default neg-full-auto** | **auto-gig** | **neg-off**

NOTE
When Gbps negotiation mode is turned off (CLI command **gig-default neg-off**), the Brocade device may inadvertently take down both ends of a link. This is a hardware limitation for which there is currently no workaround.

Port priority (QoS) modification

You can give preference to the inbound traffic on specific ports by changing the Quality of Service (QoS) level on those ports. For information and procedures, refer to “Quality of Service” chapter in the *FastIron Ethernet Switch Traffic Management Guide*.

Dynamic configuration of Voice over IP (VoIP) phones

You can configure a FastIron device to automatically detect and re-configure a VoIP phone when it is physically moved from one port to another within the same device. To do so, you must configure a **voice VLAN ID** on the port to which the VoIP phone is connected. The software stores the voice VLAN ID in the port database for retrieval by the VoIP phone.

The dynamic configuration of a VoIP phone works in conjunction with the VoiP phone discovery process. Upon installation, and sometimes periodically, a VoIP phone will query the Brocade device for VoIP information and will advertise information about itself, such as, device ID, port ID, and platform. When the Brocade device receives the VoIP phone query, it sends the voice VLAN ID in a reply packet back to the VoIP phone. The VoIP phone then configures itself within the voice VLAN.

As long as the port to which the VoIP phone is connected has a voice VLAN ID, the phone will configure itself into that voice VLAN. If you change the voice VLAN ID, the software will immediately send the new ID to the VoIP phone, and the VoIP phone will re-configure itself with the new voice VLAN.

VoIP configuration notes

- This feature works with any VoIP phone that:
  - Runs CDP
  - Sends a VoIP VLAN query message
  - Can configure its voice VLAN after receiving the VoIP VLAN reply

- Automatic configuration of a VoIP phone will not work if one of the following applies:
  - You do not configure a voice VLAN ID for a port with a VoIP phone
  - You remove the configured voice VLAN ID from a port without configuring a new one
  - You remove the port from the voice VLAN
- Make sure the port is able to intercept CDP packets (**cdp run** command).
- Some VoIP phones may require a reboot after configuring or re-configuring a voice VLAN ID. For example, if your VoIP phone queries for VLAN information only once upon boot up, you must reboot the VoIP phone before it can accept the VLAN configuration. If your phone is powered by a PoE device, you can reboot the phone by disabling then re-enabling the port.

Enabling dynamic configuration of a Voice over IP (VoIP) phone

You can create a voice VLAN ID for a port, or for a group of ports.

To create a voice VLAN ID for a port, enter commands such as the following.

```
Brocade(config)# interface ethernet 2
Brocade(config-if-e1000-2)# voice-vlan 1001
```

To create a voice VLAN ID for a group of ports, enter commands such as the following.

```
Brocade(config)# interface ethernet 1-8
Brocade(config-mif-1-8)# voice-vlan 1001
```

Syntax: [**no**] **voice-vlan** *voice-vlan-num*

where *voice-vlan-num* is a valid VLAN ID between 1 – 4095.

To remove a voice VLAN ID, use the **no** form of the command.

Viewing voice VLAN configurations

You can view the configuration of a voice VLAN for a particular port or for all ports.

To view the voice VLAN configuration for a port, specify the port number with the **show voice-vlan** command. The following example shows the command output results.

```
Brocade# show voice-vlan ethernet 2
Voice vlan ID for port 2: 1001
```

The following example shows the message that appears when the port does not have a configured voice VLAN.

```
Brocade# show voice-vlan ethernet 2
Voice vlan is not configured for port 2.
```

To view the voice VLAN for all ports, use the **show voice-vlan** command. The following example shows the command output results.

```
Brocade# show voice-vlan
Port ID    Voice-vlan
2          1001
8          150
15         200
```

Syntax: **show voice-vlan** [**ethernet** *port*]

Port flap dampening configuration

Port Flap Dampening increases the resilience and availability of the network by limiting the number of port state transitions on an interface.

If the port link state toggles from up to down for a specified number of times within a specified period, the interface is physically disabled for the specified wait period. Once the wait period expires, the port link state is re-enabled. However, if the wait period is set to zero (0) seconds, the port link state will remain disabled until it is manually re-enabled.

Port flap dampening configuration notes

- When a flap dampening port becomes a member of a trunk group, that port, as well as all other member ports of that trunk group, will inherit the primary port configuration. This means that the member ports will inherit the primary port flap dampening configuration, regardless of any previous configuration.
- The Brocade device counts the number of times a port link state toggles from "up to down", and not from "down to up".
- The sampling time or window (the time during which the specified toggle threshold can occur before the wait period is activated) is triggered when the first "up to down" transition occurs.
- "Up to down" transitions include UDLD-based toggles, as well as the physical link state.

Configuring port flap dampening on an interface

This feature is configured at the interface level.

```
Brocade(config)# interface ethernet 2/1
Brocade(config-if-e10000-2/1)# link-error-disable 10 3 10
```

Syntax: [no] link-error-disable *toggle-threshold sampling-time-in-sec wait-time-in-sec*

The *toggle-threshold* is the number of times a port link state goes from up to down and down to up before the wait period is activated. Enter a value from 1 - 50.

The *sampling-time-in-sec* is the amount of time during which the specified toggle threshold can occur before the wait period is activated. The default is 0 seconds. Enter 1 – 65535 seconds.

The *wait-time-in-sec* is the amount of time the port remains disabled (down) before it becomes enabled. Enter a value from 0 – 65535 seconds; 0 indicates that the port will stay down until an administrative override occurs.

Configuring port flap dampening on a trunk

You can configure the port flap dampening feature on the primary port of a trunk using the **link-error-disable** command. Once configured on the primary port, the feature is enabled on all ports that are members of the trunk. You cannot configure port flap dampening on port members of the trunk.

Enter commands such as the following on the primary port of a trunk.

```
Brocade(config)# interface ethernet 2/1
Brocade(config-if-e10000-2/1)# link-error-disable 10 3 10
```

Re-enabling a port disabled by port flap dampening

A port disabled by port flap dampening is automatically re-enabled once the wait period expires; however, if the wait period is set to zero (0) seconds, you must re-enable the port by entering the following command on the disabled port.

```
Brocade(config)# interface ethernet 2/1
Brocade(config-if-e10000-2/1)# no link-error-disable 10 3 10
```

Displaying ports configured with port flap dampening

Ports that have been disabled due to the port flap dampening feature are identified in the output of the **show link-error-disable** command. The following shows an example output.

```
Brocade# show link-error-disable
Port 2/1 is forced down by link-error-disable.
```

Use the **show link-error-disable all** command to display the ports with the port flap dampening feature enabled.

For FastIron Stackable devices, the output of the command shows the following.

```
Brocade# show link-error-disable all
Port8/1 is configured for link-error-disable
            threshold:1, sampling_period:10, waiting_period:0
Port8/2 is configured for link-error-disable
            threshold:1, sampling_period:10, waiting_period:0
Port8/3 is configured for link-error-disable
            threshold:1, sampling_period:10, waiting_period:0
Port8/4 is configured for link-error-disable
            threshold:1, sampling_period:10, waiting_period:0
Port8/5 is configured for link-error-disable
            threshold:4, sampling_period:10, waiting_period:2
Port8/9 is configured for link-error-disable
             threshold:2, sampling_period:20, waiting_period:0
```

For FastIron X Series devices, the output of the command shows the following.

```
Brocade# show link-error-disable all
 Port    -----------------Config---------------   ------Oper----
  #    Threshold  Sampling-Time  Shutoff-Time   State  Counter
-----   ---------  -------------  ------------   -----  -------
   11          3            120           600    Idle      N/A
   12          3            120           500    Down      424
```

Table 11 defines the port flap dampening statistics displayed by the **show link-error-disable all** command.

TABLE 11 Output of show link-error-disable

Column	Description
Port #	The port number.
Threshold	The number of times the port link state will go from up to down and down to up before the wait period is activated.
Sampling-Time	The number of seconds during which the specified toggle threshold can occur before the wait period is activated.

TABLE 11 Output of show link-error-disable (Continued)

Column	Description
Shutoff-Time	The number of seconds the port will remain disabled (down) before it becomes enabled. A zero (0) indicates that the port will stay down until an administrative override occurs.
State	The port state can be one of the following: • **Idle** – The link is normal and no link state toggles have been detected or sampled. • **Down** – The port is disabled because the number of sampled errors exceeded the configured threshold. • **Err** – The port sampled one or more errors.
Counter	• If the port state is **Idle**, this field displays **N/A**. • If the port state is **Down**, this field shows the remaining value of the shutoff timer. • If the port state is **Err**, this field shows the number of errors sampled.

Syntax: **show link-error-disable** [**all**]

Also, in FastIron X Series devices, the **show interface** command indicates if the port flap dampening feature is enabled on the port.

Example

```
Brocade# show interface ethernet 15
 GigabitEthernet15 is up, line protocol is up
 Link Error Dampening is Enabled
  Port up for 6 seconds
  Hardware is GigabitEthernet, address is 0000.0000.010e (bia 0000.0000.010e)
  Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
  Configured mdi mode AUTO, actual MDIX
Brocade# show interface ethernet 17
 GigabitEthernet17 is ERR-DISABLED, line protocol is down
  Link Error Dampening is Enabled
   Port down for 40 seconds
  Hardware is GigabitEthernet, address is 0000.0000.010e (bia 0000.0000.010e)
  Configured speed auto, actual unknown, configured duplex fdx, actual unknown
```

The line "Link Error Dampening" displays "Enabled" if port flap dampening is enabled on the port or "Disabled" if the feature is disabled on the port. The feature is enabled on the ports in the two examples above. Also, the characters "ERR-DISABLED" is displayed for the "GbpsEthernet" line if the port is disabled because of link errors.

Syntax: **show interface ethernet** *port-number*

In addition to the show commands above, the output of the **show interface brief** command for FastIron X Series indicates if a port is down due to link errors.

Example

```
Brocade# show interface brief e17

Port  Link     State      Dupl Speed Trunk Tag Priori MAC            Name
17    ERR-DIS None       None None  15    Yes level0 0000.0000.010e
```

The ERR-DIS entry under the "Link" column indicates the port is down due to link errors.

NOTE
If a port name is longer than five characters, the port name is truncated in the output of the **show interface brief** command.

Syslog messages for port flap dampening

The following Syslog messages are generated for port flap dampening.

- If the threshold for the number of times that a port link toggles from "up" to "down" then "down" to "up" has been exceeded, the following Syslog message is displayed.

  ```
  0d00h02m10s:I:ERR_DISABLE: Link flaps on port ethernet 16 exceeded threshold;
  port in err-disable state
  ```

- If the wait time (port is down) expires and the port is brought up the following Syslog message is displayed.

  ```
  0d00h02m41s:I:ERR_DISABLE: Interface ethernet 16, err-disable recovery timeout
  ```

Port loop detection

This feature allows the Brocade device to disable a port that is on the receiving end of a loop by sending test packets. You can configure the time period during which test packets are sent.

Types of loop detection

There are two types of loop detection; Strict Mode and Loose Mode. In Strict Mode, a port is disabled only if a packet is looped back to that same port. Strict Mode overcomes specific hardware issues where packets are echoed back to the input port. In Strict Mode, loop detection must be configured on the physical port.

In Loose Mode, loop detection is configured on the VLAN of the receiving port. Loose Mode disables the receiving port if packets originate from any port or VLAN on the same device. The VLAN of the receiving port must be configured for loop detection in order to disable the port.

Recovering disabled ports

Once a loop is detected on a port, it is placed in Err-Disable state. The port will remain disabled until one of the following occurs:

- You manually disable and enable the port at the Interface Level of the CLI.
- You enter the command **clear loop-detection**. This command clears loop detection statistics and enables all Err-Disabled ports.
- The device automatically re-enables the port. To set your device to automatically re-enable Err-Disabled ports, refer to "Configuring the device to automatically re-enable ports" on page 66.

Port loopback detection configuration notes

- Loopback detection packets are sent and received on both tagged and untagged ports. Therefore, this feature cannot be used to detect a loop across separate devices.

The following information applies to Loose Mode loop detection:

- With Loose Mode, two ports of a loop are disabled.
- Different VLANs may disable different ports. A disabled port affects every VLAN using it.
- Loose Mode floods test packets to the entire VLAN. This can impact system performance if too many VLANs are configured for Loose Mode loop detection.

NOTE
Brocade recommends that you limit the use of Loose Mode. If you have a large number of VLANS, configuring loop detection on all of them can significantly affect system performance because of the flooding of test packets to all configured VLANs. An alternative to configuring loop detection in a VLAN-group of many VLANs is to configure a separate VLAN with the same tagged port and configuration, and enable loop detection on this VLAN only.

NOTE
When loop detection is used with Layer 2 loop prevention protocols, such as spanning tree (STP), the Layer 2 protocol takes higher priority. Loop detection cannot send or receive probe packets if ports are blocked by Layer 2 protocols, so it does not detect Layer 2 loops when STP is running because loops within a VLAN have been prevented by STP. Loop detection running in Loose Mode can detect and break Layer 3 loops because STP cannot prevent loops across different VLANs. In these instances, the ports are not blocked and loop detection is able to send out probe packets in one VLAN and receive packets in another VLAN. In this way, loop detection running in Loose Mode disables both ingress and egress ports.

Enabling loop detection

Use the **loop-detection** command to enable loop detection on a physical port (Strict Mode) or a VLAN (Loose Mode). Loop detection is disabled by default. The following example shows a Strict Mode configuration.

```
Brocade(config)# interface ethernet 1/1
Brocade(config-if-e1000-1/1)# loop-detection
```

The following example shows a Loose Mode configuration.

```
Brocade(config)# vlan20
Brocade(config-vlan-20)# loop-detection
```

By default, the port will send test packets every one second, or the number of seconds specified by the **loop-detection-interval** command. Refer to “Configuring a global loop detection interval” on page 65.

Syntax: [**no**] **loop-detection**

Use the [no] form of the command to disable loop detection.

Configuring a global loop detection interval

The loop detection interval specifies how often a test packet is sent on a port. When loop detection is enabled, the loop detection time unit is 0.1 second, with a default of 10 (one second). The range is from 1 (one tenth of a second) to 100 (10 seconds). You can use the **show loop-detection status** command to view the loop detection interval.

To configure the global loop detection interval, enter a command similar to the following.

```
Brocade(config)# loop-detection-interval 50
```

This command sets the loop-detection interval to 5 seconds (50 x 0.1).

To revert to the default global loop detection interval of 10, enter one of the following.

```
Brocade(config)# loop-detection-interval 10
```

OR

```
Brocade(config)# no loop-detection-interval 50
```

Syntax: [**no**] **loop-detection-interval** *number*

where *number* is a value from 1 to 100. The system multiplies your entry by 0.1 to calculate the interval at which test packets will be sent.

Configuring the device to automatically re-enable ports

To configure the Brocade device to automatically re-enable ports that were disabled because of a loop detection, enter the **errdisable recovery cause loop-detection** command.

```
Brocade(config)# errdisable recovery cause loop-detection
```

The above command will cause the Brocade device to automatically re-enable ports that were disabled because of a loop detection. By default, the device will wait 300 seconds before re-enabling the ports. You can optionally change this interval to a value from 10 to 65535 seconds. Refer to "Specifying the recovery time interval" on page 66.

Syntax: [**no**] **errdisable recovery cause loop-detection**

Use the [no] form of the command to disable this feature.

Specifying the recovery time interval

The recovery time interval specifies the number of seconds the Brocade device will wait before automatically re-enabling ports that were disabled because of a loop detection. (Refer to "Configuring the device to automatically re-enable ports" on page 66.) By default, the device will wait 300 seconds. To change the recovery time interval, enter a command such as the following.

```
Brocade(config)# errdisable recovery interval 120
```

The above command configures the device to wait 120 seconds (2 minutes) before re-enabling the ports.

To revert back to the default recovery time interval of 300 seconds (5 minutes), enter one of the following commands.

```
Brocade(config)# errdisable recovery interval 300
```

OR

```
Brocade(config)# no errdisable recovery interval 120
```

Syntax: [**no**] **errdisable recovery interval** *seconds*

where *seconds* is a number from 10 to 65535.

Clearing loop-detection

To clear loop detection statistics and re-enable all ports that are in Err-Disable state because of a loop detection, enter the **clear loop-detection** command.

```
Brocade# clear loop-detection
```

Displaying loop-detection information

Use the **show loop-detection status** command to display loop detection status, as shown.

```
Brocade# show loop-detection status
loop detection packets interval: 10 (unit 0.1 sec)
Number of err-disabled ports: 3
You can re-enable err-disable ports one by one by "disable" then "enable"
under interface config, re-enable all by "clear loop-detect", or
configure "errdisable recovery cause loop-detection" for automatic recovery
index port/vlan  status                        #errdis  sent-pkts recv-pkts
1       1/13     untag, LEARNING               0         0         0
2       1/15     untag, BLOCKING               0         0         0
3       1/17     untag, DISABLED               0         0         0
4       1/18     ERR-DISABLE by itself         1         6         1
5       1/19     ERR-DISABLE by vlan 12        0         0         0
6     vlan12     2 ERR-DISABLE ports           2         24        2
```

If a port is errdisabled in Strict mode, it shows “ERR-DISABLE by itself”. If it is errdisabled due to its associated vlan, it shows “ERR-DISABLE by vlan ?”

The following command displays the current disabled ports, including the cause and the time.

```
Brocade# show loop-detection disable
Number of err-disabled ports: 3
You can re-enable err-disable ports one by one by "disable" then "enable"
under interface config, re-enable all by "clear loop-detect", or
configure "errdisable recovery cause loop-detection" for automatic recovery
index  port        caused-by    disabled-time
1      1/18        itself       00:13:30
2      1/19        vlan 12      00:13:30
3      1/20        vlan 12      00:13:30
```

This example shows the disabled ports, the cause, and the time the port was disabled. If loop-detection is configured on a physical port, the disable cause will show “itself”. For VLANs configured for loop-detection, the cause will be a VLAN.

The following command shows the hardware and software resources being used by the loop-detection feature.

```
Vlans configured loop-detection use 1 HW MAC
Vlans not configured but use HW MAC: 1 10

                      alloc in-use  avail get-fail    limit  get-mem  size init
configuration pool       16      6     10        0     3712        6    15   16
linklist pool            16     10      6        0     3712       10    16   16
```

Displaying loop detection resource information

Use the **show loop-detection resource** command to display the hardware and software resource information on loop detection.

```
Brocade# show loop-detection resource
Vlans configured loop-detection use 1 HW MAC
Vlans not configured but use HW MAC: 1 10

                      alloc in-use  avail get-fail    limit  get-mem  size init
configuration pool       16      6     10        0     3712        6    15   16
linklist pool            16     10      6        0     3712       10    16   16
```

Syntax: **show loop-detection resource**

Table 12 describes the output fields for this command.

TABLE 12 Field definitions for the show loop-detection resource command

Field	Description
This command displays the following information for the configuration pool and the linklist pool.	
alloc	Memory allocated
in-use	Memory in use
avail	Available memory
get-fail	The number of get requests that have failed
limit	The maximum memory allocation
get-mem	The number of get-memory requests
size	The size
init	The number of requests initiated

Displaying loop detection configuration status on an interface

Use the **show interface** command to display the status of loop detection configuration on a particular interface.

```
Brocade# show interface ethernet 2/1
10GigabitEthernet2/1 is up, line protocol is up
Port up for 1 day 22 hours 43 minutes 5 seconds
Hardware is 10GigabitEthernet, address is 0000.0089.1100 (bia 0000.0089.1118)
Configured speed 10Gbit, actual 10Gbit, configured duplex fdx, actual fdx
Member of 9 L2 VLANs, port is tagged, port state is FORWARDING
BPDU guard is Disabled, ROOT protect is Disabled
Link Error Dampening is Disabled
STP configured to ON, priority is level0
Loop Detection is ENABLED
Flow Control is enabled
Mirror disabled, Monitor disabled
Member of active trunk ports 2/1,2/2, primary port
Member of configured trunk ports 2/1,2/2, primary port
No port name
IPG XGMII 96 bits-time
MTU 1500 bytes, encapsulation ethernet
ICL port for BH1 in cluster id 1
300 second input rate: 2064 bits/sec, 3 packets/sec, 0.00% utilization
300 second output rate: 768 bits/sec, 1 packets/sec, 0.00% utilization
```

```
171319 packets input, 12272674 bytes, 0 no buffer
Received 0 broadcasts, 63650 multicasts, 107669 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants
51094 packets output, 3925313 bytes, 0 underruns
Transmitted 2 broadcasts, 42830 multicasts, 8262 unicasts
0 output errors, 0 collisions
Relay Agent Information option: Disabled
```

Syslog message due to disabled port in loop detection

The following message is logged when a port is disabled due to loop detection. This message also appears on the console.

```
loop-detect: port ?\?\? vlan ?, into errdisable state
```

The Errdisable function logs a message whenever it re-enables a port.

Chapter

Operations, Administration, and Maintenance 3

Table 13 lists the individual Brocade FastIron switches and the operations, administration, and maintenance (OAM) features they support. These features are supported in the Layer 2 and Layer 3 software images, except where explicitly noted.

TABLE 13 Supported operations, administration, and maintenance features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6430	ICX 6450
Flash and boot code verification	Yes	Yes	Yes	Yes	Yes
Flash image verification	Yes	Yes	Yes	Yes	Yes
Software upgrade via CLI	Yes	Yes	Yes	Yes	Yes
Software upgrade via SNMP	Yes	Yes	Yes	Yes	Yes
Hitless management: • Hitless switchover • Hitless failover • Hitless OS upgrade	Yes (Hitless switchover and Hitless failover only; Hitless OS upgrade is not supported) Refer to “Hitless stacking” section in the *FastIron Ethernet Switch Stacking Configuration Guide*				
Hitless support: • PBR • GRE Tunnels • IPv6 to IPv4 Tunnels	Yes	Yes (PBR only)	Yes (PBR only)	No	No
Boot code synchronization for active and redundant management modules	Yes	No	No	No	No
Software reboot	Yes	Yes	Yes	Yes	Yes
Show boot preference	Yes	Yes	Yes	Yes	Yes
Load and save configuration files	Yes	Yes	Yes	Yes	Yes
System reload scheduling	Yes	Yes	Yes	Yes	Yes
Diagnostic error codes and remedies for TFTP transfers	Yes	Yes	Yes	Yes	Yes
IPv4 ping	Yes	Yes	Yes	Yes	Yes
IPv4 traceroute	Yes	Yes	Yes	Yes	Yes
Layer 3 hitless route purge	Yes [1]	Yes	Yes	No	Yes

1. Third generation modules.

This chapter contains the following sections:

OAM Overview

For easy software image management, all Brocade devices support the download and upload of software images between the flash modules on the devices and a Trivial File Transfer Protocol (TFTP) server on the network.

Brocade devices have two flash memory modules:

- *Primary flash* – The default local storage device for image files and configuration files.
- *Secondary flash* – A second flash storage device. You can use the secondary flash to store redundant images for additional booting reliability or to preserve one software image while testing another one.

Only one flash device is active at a time. By default, the primary image will become active upon reload.

You can update the software contained on a flash module using TFTP to copy the update image from a TFTP server onto the flash module. In addition, you can copy software images and configuration files from a flash module to a TFTP server.

NOTE
Brocade devices are TFTP clients but not TFTP servers. You must perform the TFTP transaction from the Brocade device. You cannot "put" a file onto the Brocade device using the interface of your TFTP server.

NOTE
If you are attempting to transfer a file using TFTP but have received an error message, refer to "Diagnostic error codes and remedies for TFTP transfers" on page 94.

Software versions installed and running on a device

Use the following methods to display the software versions running on the device and the versions installed in flash memory.

Determining the flash image version running on the device

To determine the flash image version running on a device, enter the **show version** command at any level of the CLI. Some examples are shown below.

Compact devices

To determine the flash image version running on a Compact device, enter the **show version** command at any level of the CLI. The following shows an example output.

```
Brocade#show version
Copyright (c) 1996-2012 Brocade Communications Systems, Inc. All rights reserved.
    UNIT 1: compiled on Mar  2 2012 at 12:38:17 labeled as ICX64S07400
                (10360844 bytes) from Primary ICX64S07400.bin
        SW: Version 07.4.00T311
  Boot-Monitor Image size = 774980, Version:07.4.00T310 (kxz07400)
  HW: Stackable ICX6450-24
==========================================================================
UNIT 1: SL 1: ICX6450-24 24-port Management Module
         Serial  #: BZSxxxxxxxx
         License: BASE_SOFT_PACKAGE   (LID: dbuFJJHiFFi)
         P-ENGINE  0: type DEF0, rev 01
==========================================================================
UNIT 1: SL 2: ICX6450-SFP-Plus 4port 40G Module
==========================================================================
  800 MHz ARM processor ARMv5TE, 400 MHz bus
65536 KB flash memory
  512 MB DRAM
STACKID 1  system uptime is 3 minutes 39 seconds
The system : started=warm start  reloaded=by "reload"
```

The version information is shown in bold type in this example:

- "03.0.00T53" indicates the flash code version number. The "T53" is used by Brocade for record keeping.
- "labeled as FER03000" indicates the flash code image label. The label indicates the image type and version and is especially useful if you change the image file name.
- "Primary fer03000.bin" indicates the flash code image file name that was loaded.

Displaying flash image version on chassis devices

To determine the flash image version running on a chassis device, enter the **show version** command at any level of the CLI. The following is an example output.

```
Brocade#show version
==========================================================================
Active Management CPU [Slot-9]:
  SW: Version 07.4.00T3e3 Copyright (c) 1996-2012 Brocade Communications Systems,
Inc. All rights reserved.
      Compiled on Mar 02 2012 at 11:54:29 labeled as SXR07400
      (4585331 bytes) Primary /GA/SXR07400.bin
      BootROM: Version 07.2.00T3e5 (FEv2)
      Chassis Serial  #: Bxxxxxxxxx
         License: SX_V6_HW_ROUTER_IPv6_SOFT_PACKAGE   (LID: yGFJGOiFLd)
  HW: Chassis FastIron SX 800-PREM6 (PROM-TYPE SX-FIL3U-6-IPV6)
==========================================================================
Standby Management CPU [Slot-10]:
  SW: Version 07.4.00T3e3 Copyright (c) 1996-2012 Brocade Communications Systems,
Inc. All rights reserved.
      Compiled on Mar 02 2012 at 11:54:29 labeled as SXR07400
      BootROM: Version 07.2.00T3e5 (FEv2)
  HW: Chassis FastIron SX 800-PREM6 (PROM-TYPE SX-FIL3U-6-IPV6)
==========================================================================
SL 1: SX-FI-8XG 8-port 10G Fiber
```

```
      Serial  #: BQKxxxxxxxx
      P-ASIC  0: type C341, rev 00  subrev 00
==========================================================================
SL 2: SX-FI-24GPP 24-port Gig Copper + PoE+
      Serial  #: BTUxxxxxxxx
      P-ASIC  2: type C300, rev 00  subrev 00
==========================================================================
SL 8: SX-FI-48GPP 48-port Gig Copper + PoE+
      Serial  #: BFVxxxxxxxx
      P-ASIC 14: type C300, rev 00  subrev 00
==========================================================================
SL 9: SX-FIZMR6 0-port Management
      Serial  #: Wxxxxxxxxx
         License: SX_V6_HW_ROUTER_IPv6_SOFT_PACKAGE   (LID: yGFJGOiFLd)
==========================================================================
SL 10: SX-FIZMR6 0-port Management
      Serial  #: Wxxxxxxxxx
         License: SX_V6_HW_ROUTER_IPv6_SOFT_PACKAGE   (LID: ÿÿÿÿÿÿÿÿÿÿ)
==========================================================================
Active Management Module:
  660 MHz Power PC processor 8541 (version 0020/0020) 66 MHz bus
  512 KB boot flash memory
16384 KB code flash memory
  512 MB DRAM
Standby Management Module:
  660 MHz Power PC processor 8541 (version 0020/0020) 66 MHz bus
  512 KB boot flash memory
16384 KB code flash memory
  512 MB DRAM
The system uptime is 1 minutes 2 seconds
The system : started=warm start   reloaded=by "reload"
```

The version information is shown in bold type in this example:

- "03.1.00aT3e3" indicates the flash code version number. The "T3e3" is used by Brocade for record keeping.
- "labeled as SXR03100a" indicates the flash code image label. The label indicates the image type and version and is especially useful if you change the image file name.
- "Primary SXR03100a.bin" indicates the flash code image file name that was loaded.

Displaying the boot image version running on the device

To determine the boot image running on a device, enter the **show flash** command at any level of the CLI. The following shows an example output.

```
Brocade#show flash
Active Management Module (Slot 9):
Compressed Pri Code size = 3613675, Version 03.1.00aT3e3 (sxr03100a.bin)
Compressed Sec Code size = 2250218, Version 03.1.00aT3e1 (sxs03100a.bin)
Compressed BootROM Code size = 524288, Version 03.0.01T3e5
Code Flash Free Space = 9699328
Standby Management Module (Slot 10):
Compressed Pri Code size = 3613675, Version 03.1.00aT3e3 (sxr03100a.bin)
Compressed Sec Code size = 2250218, Version 03.1.00aT3e1 (sxs03100a.bin)
Compressed BootROM Code size = 524288, Version 03.0.01T3e5
Code Flash Free Space = 524288
```

The boot code version is shown in bold type.

Displaying the image versions installed in flash memory

Enter the **show flash** command to display the boot and flash images installed on the device. An example of the command output is shown in "Displaying the boot image version running on the device" on page 74:

- The "Compressed Pri Code size" line lists the flash code version installed in the primary flash area.
- The "Compressed Sec Code size" line lists the flash code version installed in the secondary flash area.
- The "Boot Monitor Image size" line lists the boot code version installed in flash memory. The device does not have separate primary and secondary flash areas for the boot image. The flash memory module contains only one boot image.

NOTE
To minimize the boot-monitor image size on FastIron devices, the **ping** and **tftp** operations performed in the boot-monitor mode are restricted to copper ports on the FastIron Chassis management modules and to copper ports on the FastIron stackable switch combination copper and fiber ports. The fiber ports on these devices do not have the ability to **ping** or **tftp** from the boot-monitor mode.

Flash image verification

The Flash Image Verification feature allows you to verify boot images based on hash codes, and to generate hash codes where needed. This feature lets you select from three data integrity verification algorithms:

- **MD5** - Message Digest algorithm (RFC 1321)
- **SHA1** - US Secure Hash Algorithm (RFC 3174)
- **CRC** - Cyclic Redundancy Checksum algorithm

Flash image CLI commands

Use the following command syntax to verify the flash image:

Syntax: verify md5 | **sha1** | **crc32** *ASCII string* | **primary** | **secondary** [*hash code*]

- **md5** – Generates a 16-byte hash code
- **sha1** – Generates a 20-byte hash code
- **crc32** – Generates a 4 byte checksum
- **ascii string** – A valid image filename
- **primary** – The primary boot image (primary.img)
- **secondary** – The secondary boot image (secondary.img)
- **hash code** – The hash code to verify

The following examples show how the **verify** command can be used in a variety of circumstances.

To generate an MD5 hash value for the secondary image, enter the following command.

```
Brocade#verify md5 secondary
Brocade#........................Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653862
```

To generate a SHA-1 hash value for the secondary image, enter the following command.

```
Brocade#verify sha secondary
Brocade#........................Done
Size = 2044830, SHA1 49d12d26552072337f7f5fcaef4cf4b742a9f525
```

To generate a CRC32 hash value for the secondary image, enter the following command.

```
Brocade#verify crc32 secondary
Brocade#........................Done
Size = 2044830, CRC32 b31fcbc0
```

To verify the hash value of a secondary image with a known value, enter the following commands.

```
Brocade#verify md5 secondary 01c410d6d153189a4a5d36c955653861
Brocade#........................Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653862
Verification FAILED.
```

In the previous example, the codes did not match, and verification failed. If verification succeeds, the output will look like this.

```
Brocade#verify md5 secondary 01c410d6d153189a4a5d36c955653861
Brocade#........................Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653861
Verification SUCEEDED.
```

The following examples show this process for SHA-1 and CRC32 algorithms.

```
Brocade#verify sha secondary 49d12d26552072337f7f5fcaef4cf4b742a9f525
Brocade#........................Done
Size = 2044830, sha 49d12d26552072337f7f5fcaef4cf4b742a9f525
Verification SUCCEEDED.
```

and

```
Brocade#verify crc32 secondary b31fcbc0
Brocade#........................Done
Size = 2044830, CRC32 b31fcbc0
Verification SUCCEEDED.
```

Image file types

This section lists the boot and flash image file types supported and how to install them on the FastIron family of switches. For information about a specific version of code, refer to the release notes.

TABLE 14 Software image files

Product	Boot image[1]	Flash image
FSX 800 FSX 1600	sxzxxxxx.bin	SXLSxxxxx.bin (Layer 2) or SXLRxxxxx.bin (full Layer 3)

TABLE 14 Software image files (Continued)

Product	Boot image[1]	Flash image
FCX ICX 6610	grzxxxxxx.bin	FCXSxxxxx.bin (Layer 2) or FCXRxxxxx.bin (Layer 3)
ICX 6430 ICX 6450	kxzxxxxx.bin	ICX64Sxxxxx.bin (Layer 2) or ICX64Rxxxxx.bin (Layer 3 - ICX 6450 only)

1. These images are applicable to these devices only and are not interchangeable. For example, you cannot load FCX boot or flash images on a FSX device, and vice versa.

Software upgrades

For instructions about upgrading the software, refer to *FastIron Ethernet Switch Software Upgrade Guide*.

Boot code synchronization feature

The Brocade device supports automatic synchronization of the boot image in the active and redundant management modules. When the new boot image is copied into the active module, it is automatically synchronized with the redundant management module.

NOTE
There is currently no option for manual synchronization of the boot image.

To activate the boot synchronization process, enter the following command.

```
Brocade#copy tftp flash 10.20.65.194 /GA/SXZ07200.bin bootrom
```

The system responds with the following message.

```
Brocade#Load to buffer (8192 bytes per dot)
..................Write to boot flash......................
TFTP to Flash Done.
Brocade#Synchronizing with standby module...
Boot image synchronization done.
```

Viewing the contents of flash files

The **copy flash console** command can be used to display the contents of a configuration file, backup file, or renamed file stored in flash memory. The file contents are displayed on the console when the command is entered at the CLI.

To display a list of files stored in flash memory, do one of the following:

- For devices other than FCX and ICX, enter the **dir** command at the monitor mode. To enter monitor mode from any level of the CLI, press the **Shift** and **Control+Y** keys simultaneously then press the M key. Enter the **dir** command to display a list of the files stored in flash memory. To exit monitor mode and return to the CLI, press **Control+Z**.

- For FCX devices, enter the **show dir** command at any level of the CLI, or enter the **dir** command at the monitor mode.
- For ICX devices, enter the **show files** command at the device configuration prompt.

The following shows an example command output.

```
Brocade#show dir
133 [38f4] boot-parameter
         0 [ffff] bootrom
   3802772 [0000] primary
   4867691 [0000] secondary
       163 [dd8e] stacking.boot
      1773 [0d2d] startup-config
      1808 [acfa] startup-config.backup
   8674340 bytes 7 File(s)
  56492032 bytes free
```

Syntax: **show dir**

To display the contents of a flash configuration file, enter a command such as the following from the User EXEC or Privileged EXEC mode of the CLI:

```
Brocade#copy flash console startup-config.backup
ver 07.0.00b1T7f1 !
stack unit 1
  module 1 fcx-24-port-management-module
  module 2 fcx-cx4-2-port-16g-module
  module 3 fcx-xfp-2-port-10g-module
  priority 80
  stack-port 1/2/1 1/2/2
stack unit 2
  module 1 fcx-48-poe-port-management-module
  module 2 fcx-cx4-2-port-16g-module
  module 3 fcx-xfp-2-port-10g-module
  stack-port 2/2/1 2/2/2
stack enable
!
!
!
!
vlan 1 name DEFAULT-VLAN by port
 no spanning-tree
 metro-rings 1
 metro-ring 1
  master
  ring-interfaces  ethernet 1/1/2  ethernet 1/1/3
  enable
!
vlan 10 by port
 mac-vlan-permit ethe 1/1/5 to 1/1/6 ethe 2/1/5 to 2/1/6  no spanning-tree !
vlan 20 by port
 untagged ethe 1/1/7 to 1/1/8
 no spanning-tree
 pvlan type primary
 pvlan mapping 40 ethe 1/1/8
 pvlan mapping 30 ethe 1/1/7
!
vlan 30 by port
 untagged ethe 1/1/9 to 1/1/10
 no spanning-tree
 pvlan type community
!
...
some lines omitted for brevity...
```

Syntax: **copy flash console** *filename*

For *filename*, enter the name of a file stored in flash memory.

Using SNMP to upgrade software

You can use a third-party SNMP management application such as HP OpenView to upgrade software on a Brocade device.

NOTE
The syntax shown in this section assumes that you have installed HP OpenView in the "/usr" directory.

NOTE
Brocade recommends that you make a backup copy of the startup-config file before you upgrade the software. If you need to run an older release, you will need to use the backup copy of the startup-config file.

---

1. Configure a read-write community string on the Brocade device, if one is not already configured. To configure a read-write community string, enter the following command from the global CONFIG level of the CLI.

   snmp-server community *string* **ro | rw**

   where *string* is the community string and can be up to 32 characters long.

2. On the Brocade device, enter the following command from the global CONFIG level of the CLI.

   no snmp-server pw-check

   This command disables password checking for SNMP set requests. If a third-party SNMP management application does not add a password to the password field when it sends SNMP set requests to a Brocade device, by default the Brocade device rejects the request.

3. From the command prompt in the UNIX shell, enter the following command.

   /usr/OV/bin/snmpset -c *rw-community-string brcd-ip-addr* **1.3.6.1.4.1.1991.1.1.2.1.5.0 ipaddress** *tftp-ip-addr* **1.3.6.1.4.1.1991.1.1.2.1.6.0 octetstringascii** *file-name* **1.3.6.1.4.1.1991.1.1.2.1.7.0 integer** *command-integer*

   where

   rw-community-string is a read-write community string configured on the Brocade device.

   brcd-ip-addr is the IP address of the Brocade device.

   tftp-ip-addr is the TFTP server IP address.

   file-name is the image file name.

   command-integer is one of the following.

   20 – Download the flash code into the primary flash area.

   22 – Download the flash code into the secondary flash area.

Software reboot

You can use boot commands to immediately initiate software boots from a software image stored in primary or secondary flash on a Brocade device or from a BootP or TFTP server. You can test new versions of code on a Brocade device or choose the preferred boot source from the console boot prompt without requiring a system reset.

NOTE
It is very important that you verify a successful TFTP transfer of the boot code *before* you reset the system. If the boot code is not transferred successfully but you try to reset the system, the system will not have the boot code with which to successfully boot.

By default, the Brocade device first attempts to boot from the image stored in its primary flash, then its secondary flash, and then from a TFTP server. You can modify this booting sequence at the global CONFIG level of the CLI using the **boot system...** command.

NOTE
FSX device with FastIron 08.0.00a, ICX 6430, and ICX 6450 devices support only one configured system boot preference.

To initiate an immediate boot from the CLI, enter one of the **boot system...** commands.

NOTE
When using the **boot system tftp** command, the IP address of the device and the TFTP server should be in the same subnet.

Software boot configuration notes

- In FastIron X Series devices, the **boot system tftp** command is supported on ports e 1 through e 12 only.
- If you are booting the device from a TFTP server through a fiber connection, use the following command: **boot system tftp** *ip-address filename* **fiber-port**.
- The **boot system tftp** command is not supported in a stacking environment.

Displaying the boot preference

Use the **show boot-preference** command to display the boot sequence in the startup config and running config files. The boot sequence displayed is also identified as either user-configured or the default.

The following example shows the default boot sequence preference.

```
Brocade#show boot-preference
Boot system preference (Configured):
   Use Default
Boot system preference(Default):
   Boot system flash primary
   Boot system flash secondary
```

The following example shows a user-configured boot sequence preference.

```
Brocade#show boot-preference
Boot system preference(Configured):
        Boot system tftp 10.1.1.1 FCXR08000.bin
        Boot system flash primary

Boot system preference(Default):
        Boot system flash primary
        Boot system flash secondary
```

Syntax: show boot-preference

The results of the **show run** command for the configured example above appear as follows.

```
Brocade#show run
Current configuration:
!
ver 08.0.00T7f3
!
stack unit 1
  module 1 fcx-24-poe-port-management-module
  module 2 fcx-cx4-2-port-16g-module
  priority 128
  stack-port 1/2/1 1/2/2
stack unit 2
  module 1 fcx-48-port-management-module
  module 2 fcx-cx4-2-port-16g-module
  stack-port 2/2/1 2/2/2
stack enable
stack mac 748e.f80e.dcc0
!
boot sys tf 10.1.1.1 FCXR08000.bin
boot sys fl pri
ip route 0.0.0.0/0 10.37.234.129
!
end
```

Loading and saving configuration files

For easy configuration management, all Brocade devices support both the download and upload of configuration files between the devices and a TFTP server on the network.

You can upload either the startup configuration file or the running configuration file to the TFTP server for backup and use in booting the system:

- *Startup configuration file* – This file contains the configuration information that is currently saved in flash. To display this file, enter the **show configuration** command at any CLI prompt.
- *Running configuration file* – This file contains the configuration active in the system RAM but not yet saved to flash. These changes could represent a short-term requirement or general configuration change. To display this file, enter the **show running-config** or **write terminal** command at any CLI prompt.

Each device can have one startup configuration file and one running configuration file. The startup configuration file is shared by both flash modules. The running configuration file resides in DRAM.

When you load the startup-config file, the CLI parses the file three times.

1. During the first pass, the parser searches for **system-max** commands. A **system-max** command changes the size of statically configured memory.
2. During the second pass, the parser implements the **system-max** commands if present and also implements trunk configuration commands (**trunk** command) if present.
3. During the third pass, the parser implements the remaining commands.

Replacing the startup configuration with the running configuration

After you make configuration changes to the active system, you can save those changes by writing them to flash memory. When you write configuration changes to flash memory, you replace the startup configuration with the running configuration.

To replace the startup configuration with the running configuration, enter the following command at any Enable or CONFIG command prompt.

```
Brocade#write memory
```

Replacing the running configuration with the startup configuration

If you want to back out of the changes you have made to the running configuration and return to the startup configuration, enter the following command at the Privileged EXEC level of the CLI.

```
Brocade#reload
```

Logging changes to the startup-config file

You can configure a Brocade device to generate a Syslog message when the startup-config file is changed. The trap is enabled by default.

The following Syslog message is generated when the startup-config file is changed.

```
startup-config was changed
```

If the startup-config file was modified by a valid user, the following Syslog message is generated.

```
startup-config was changed by username
```

To disable or re-enable Syslog messages when the startup-config file is changed, use the following command.

Syntax: [no] **logging enable config-changed**

Copying a configuration file to or from a TFTP server

To copy the startup-config or running-config file to or from a TFTP server, use one of the following methods.

NOTE
For details about the **copy** and **ncopy** commands used with IPv6, refer to "Using the IPv6 copy command" on page 88and "IPv6 ncopy command" on page 90.

NOTE
You can name the configuration file when you copy it to a TFTP server. However, when you copy a configuration file from the server to a Brocade device, the file is always copied as "startup-config" or "running-config", depending on which type of file you saved to the server.

To initiate transfers of configuration files to or from a TFTP server using the CLI, enter one of the following commands:

- **copy startup-config tftp** *tftp-ip-addr filename* – Use this command to upload a copy of the startup configuration file from the Layer 2 Switch or Layer 3 Switch to a TFTP server.
- **copy running-config tftp** *tftp-ip-addr filename* – Use this command to upload a copy of the running configuration file from the Layer 2 Switch or Layer 3 Switch to a TFTP server.
- **copy tftp startup-config** *tftp-ip-addr filename* – Use this command to download a copy of the startup configuration file from a TFTP server to a Layer 2 Switch or Layer 3 Switch.

NOTE
It is recommended to use a script or the **copy running-config tftp** command for extensive configuration. You should not copy-paste configuration with more than 2000 characters into CLI.

Dynamic configuration loading

You can load dynamic configuration commands (commands that do not require a reload to take effect) from a file on a TFTP server into the running-config on the Brocade device. You can make configuration changes off-line, then load the changes directly into the device running-config, without reloading the software.

Dynamic configuration usage considerations

- Use this feature only to load configuration information that does not require a software reload to take effect. For example, you cannot use this feature to change statically configured memory (**system-max** command) or to enter trunk group configuration information into the running-config.
- Do not use this feature if you have deleted a trunk group but have not yet placed the changes into effect by saving the configuration and then reloading. When you delete a trunk group, the command to configure the trunk group is removed from the device running-config, but the trunk group remains active. To finish deleting a trunk group, save the configuration (to the startup-config file), then reload the software. After you reload the software, then you can load the configuration from the file.
- Do not load port configuration information for secondary ports in a trunk group. Since all ports in a trunk group use the port configuration settings of the primary port in the group, the software cannot implement the changes to the secondary port.

Preparing the configuration file

A configuration file that you create must follow the same syntax rules as the startup-config file the device creates.

- The configuration file is a script containing CLI configuration commands. The CLI reacts to each command entered from the file in the same way the CLI reacts to the command if you enter it. For example, if the command results in an error message or a change to the CLI configuration level, the software responds by displaying the message or changing the CLI level.
- The software retains the running-config that is currently on the device, and changes the running-config only by adding new commands from the configuration file. If the running config already contains a command that is also in the configuration file you are loading, the CLI rejects the new command as a duplicate and displays an error message. For example, if the running-config already contains a a command that configures ACL 1, the software rejects ACL 1 in the configuration file, and displays a message that ACL 1 is already configured.

- The file can contain global CONFIG commands or configuration commands for interfaces, routing protocols, and so on. You cannot enter User EXEC or Privileged EXEC commands.
- The default CLI configuration level in a configuration file is the global CONFIG level. Thus, the first command in the file must be a global CONFIG command or " ! ". The ! (exclamation point) character means "return to the global CONFIG level".

NOTE
You can enter text following " ! " as a comment. However, the " !" is not a comment marker. It returns the CLI to the global configuration level.

NOTE
If you copy-and-paste a configuration into a management session, the CLI ignores the " ! " instead of changing the CLI to the global CONFIG level. As a result, you might get different results if you copy-and-paste a configuration instead of loading the configuration using TFTP.

- Make sure you enter each command at the correct CLI level. Since some commands have identical forms at both the global CONFIG level and individual configuration levels, if the CLI response to the configuration file results in the CLI entering a configuration level you did not intend, then you can get unexpected results.

  For example, if a trunk group is active on the device, and the configuration file contains a command to disable STP on one of the secondary ports in the trunk group, the CLI rejects the commands to enter the interface configuration level for the port and moves on to the next command in the file you are loading. If the next command is a spanning-tree command whose syntax is valid at the global CONFIG level as well as the interface configuration level, then the software applies the command globally. Here is an example.

  The configuration file contains these commands.

```
interface ethernet 2
no spanning-tree
```

  The CLI responds like this.

```
Brocade(config)#interface ethernet 2
Error - cannot configure secondary ports of a trunk
Brocade(config)#no spanning-tree
Brocade(config)#
```

- If the file contains commands that must be entered in a specific order, the commands must appear in the file in the required order. For example, if you want to use the file to replace an IP address on an interface, you must first remove the old address using "no" in front of the **ip address** command, then add the new address. Otherwise, the CLI displays an error message and does not implement the command. Here is an example.

  The configuration file contains these commands.

```
interface ethernet 11
ip address 10.10.10.69/24
```

  The running-config already has a command to add an address to port 11, so the CLI responds like this.

```
Brocade(config)#interface ethernet 11
Brocade(config-if-e1000-11)#ip add 10.10.10.69/24
Error: can only assign one primary ip address per subnet
Brocade(config-if-e1000-11)#
```

  To successfully replace the address, enter commands into the file as follows.

```
interface ethernet 11
no ip address 10.20.20.69/24
ip address 10.10.10.69/24
```

This time, the CLI accepts the command, and no error message is displayed.

```
Brocade(config)#interface ethernet 11
Brocade(config-if-e1000-11)#no ip add 10.20.20.69/24
Brocade(config-if-e1000-111)#ip add 10.10.10.69/24
Brocade(config-if-e1000-11)
```

- Always use the **end** command at the end of the file. The **end** command must appear on the last line of the file, by itself.

Loading the configuration information into the running-config

To load the file from a TFTP server, use either of the following commands:

- **copy tftp running-config** *ip-addr filename*
- **ncopy tftp** *ip-addr filename* **running-config**

NOTE
In FastIron 08.0.00a, the **copy tftp running-config** command merges only the access-lists and mac-filters configuration from the configuration file on the TFTP server to the running configuration on the device.

NOTE
If you are loading a configuration file that uses a truncated form of the CLI command **access-list**, the software will not go into batch mode.

For example, the following command line *will initiate* batch mode.

```
access-list 131 permit host pc1 host pc2
```

The following command line *will not* initiate batch mode.

```
acc 131 permit host pc1 host pc2
```

Maximum file sizes for startup-config file and running-config

Each Brocade device has a maximum allowable size for the running-config and the startup-config file. If you use TFTP to load additional information into a device running-config or startup-config file, it is possible to exceed the maximum allowable size. If this occurs, you will not be able to save the configuration changes.

The maximum size for the running-config and the startup-config file is 640K each.

To determine the size of a running-config or startup-config file, copy it to a TFTP server, then use the directory services on the server to list the size of the copied file. To copy the running-config or startup-config file to a TFTP server, use one of the following commands:

- Commands to copy the running-config to a TFTP server:
  - **copy running-config tftp** *ip-addr filename*
  - **ncopy running-config tftp** *ip-addr from-name*
- Commands to copy the startup-config file to a TFTP server:
  - **copy startup-config tftp** *ip-addr filename*

- **ncopy startup-config tftp** *ip-addr from-name*

Loading and saving configuration files with IPv6

This section describes the IPv6 **copy** and **ncopy** commands.

Using the IPv6 copy command

The **copy** command for IPv6 allows you to do the following:

- Copy a file from a specified source to an IPv6 TFTP server
- Copy a file from an IPv6 TFTP server to a specified destination

Copying a file to an IPv6 TFTP server

You can copy a file from the following sources to an IPv6 TFTP server:

- Flash memory
- Running configuration
- Startup configuration

Copying a file from flash memory

For example, to copy the primary or secondary boot image from the device flash memory to an IPv6 TFTP server, enter a command such as the following.

```
Brocade#copy flash tftp 2001:DB8:e0ff:7837::3 test.img secondary
```

This command copies the secondary boot image named test.img from flash memory to a TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3.

Syntax: **copy flash tftp** *ipv6-address source-file-name* **primary** | **secondary**

The *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *source-file-name* parameter specifies the name of the file you want to copy to the IPv6 TFTP server.

The **primary** keyword specifies the primary boot image, while the **secondary** keyword specifies the secondary boot image.

Copying a file from the running or startup configuration

For example, to copy the running configuration to an IPv6 TFTP server, enter a command such as the following.

```
Brocade#copy running-config tftp 2001:DB8:e0ff:7837::3 newrun.cfg
```

This command copies the running configuration to a TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3 and names the file on the TFTP server newrun.cfg.

Syntax: **copy running-config** | **startup-config tftp** *ipv6-address destination-file-name*

Specify the **running-config** keyword to copy the running configuration file to the specified IPv6 TFTP server.

Specify the **startup-config** keyword to copy the startup configuration file to the specified IPv6 TFTP server.

The tftp *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *destination-file-name* parameter specifies the name of the file that is copied to the IPv6 TFTP server.

Copying a file from an IPv6 TFTP server

You can copy a file from an IPv6 TFTP server to the following destinations:

- Flash memory
- Running configuration
- Startup configuration

Copying a file to flash memory

For example, to copy a boot image from an IPv6 TFTP server to the primary or secondary storage location in the device flash memory, enter a command such as the following.

```
Brocade#copy tftp flash 2001:DB8:e0ff:7837::3 test.img secondary
```

This command copies a boot image named test.img from an IPv6 TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3 to the secondary storage location in the device flash memory.

Syntax: **copy tftp flash** *ipv6-address source-file-name* **primary** | **secondary**

The *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *source-file-name* parameter specifies the name of the file you want to copy from the IPv6 TFTP server.

The **primary** keyword specifies the primary storage location in the device flash memory, while the **secondary** keyword specifies the secondary storage location in the device flash memory.

Copying a file to the running or startup configuration

For example, to copy a configuration file from an IPv6 TFTP server to the running or startup configuration, enter a command such as the following.

```
Brocade#copy tftp running-config 2001:DB8:e0ff:7837::3 newrun.cfg overwrite
```

This command copies the newrun.cfg file from the IPv6 TFTP server and overwrites the running configuration file with the contents of newrun.cfg.

NOTE
To activate this configuration, you must reload (reset) the device.

Syntax: **copy tftp running-config** | **startup-config** *ipv6-address source-file-name* [**overwrite**]

Specify the **running-config** keyword to copy the running configuration from the specified IPv6 TFTP server.

The *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *source-file-name* parameter specifies the name of the file that is copied from the IPv6 TFTP server.

The **overwrite** keyword specifies that the device should overwrite the current configuration file with the copied file. If you do not specify this parameter, the device copies the file into the current running or startup configuration but does not overwrite the current configuration.

IPv6 ncopy command

The **ncopy** command for IPv6 allows you to do the following:

- Copy a primary or secondary boot image from flash memory to an IPv6 TFTP server.
- Copy the running configuration to an IPv6 TFTP server.
- Copy the startup configuration to an IPv6 TFTP server
- Upload various files from an IPv6 TFTP server.

Copying a primary or secondary boot Image from flash memory to an IPv6 TFTP server

For example, to copy the primary or secondary boot image from the device flash memory to an IPv6 TFTP server, enter a command such as the following.

```
Brocade#ncopy flash primary tftp 2001:DB8:e0ff:7837::3 primary.img
```

This command copies the primary boot image named primary.img from flash memory to a TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3.

Syntax: **ncopy flash primary** | **secondary tftp** *ipv6-address source-file-name*

The **primary** keyword specifies the primary boot image, while the **secondary** keyword specifies the secondary boot image.

The **tftp** *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *source-file-name* parameter specifies the name of the file you want to copy from flash memory.

Copying the running or startup configuration to an IPv6 TFTP server

For example, to copy a device running or startup configuration to an IPv6 TFTP server, enter a command such as the following.

```
Brocade#ncopy running-config tftp 2001:DB8:e0ff:7837::3 bakrun.cfg
```

This command copies a device running configuration to a TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3 and names the destination file bakrun.cfg.

Syntax: **ncopy running-config** | **startup-config tftp** *ipv6-address destination-file-name*

Specify the **running-config** keyword to copy the device running configuration or the **startup-config** keyword to copy the device startup configuration.

The **tftp** *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *destination-file-name* parameter specifies the name of the running configuration that is copied to the IPv6 TFTP server.

IPv6 TFTP server file upload

You can upload the following files from an IPv6 TFTP server:

- Primary boot image.
- Secondary boot image.
- Running configuration.
- Startup configuration.

Uploading a primary or secondary boot image from an IPv6 TFTP server

For example, to upload a primary or secondary boot image from an IPv6 TFTP server to a device flash memory, enter a command such as the following.

```
Brocade#ncopy tftp 2001:DB8:e0ff:7837::3 primary.img flash primary
```

This command uploads the primary boot image named primary.img from a TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3 to the device primary storage location in flash memory.

Syntax: **ncopy tftp** *ipv6-address source-file-name* **flash primary** | **secondary**

The **tftp** *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *source-file-name* parameter specifies the name of the file you want to copy from the TFTP server.

The **primary** keyword specifies the primary location in flash memory, while the **secondary** keyword specifies the secondary location in flash memory.

Uploading a running or startup configuration from an IPv6 TFTP server

For example to upload a running or startup configuration from an IPv6 TFTP server to a device, enter a command such as the following.

```
Brocade#ncopy tftp 2001:DB8:e0ff:7837::3 newrun.cfg running-config
```

This command uploads a file named newrun.cfg from a TFTP server with the IPv6 address of 2001:DB8:e0ff:7837::3 to the device.

Syntax: **ncopy tftp** *ipv6-address source-file-name* **running-config** | **startup-config**

The **tftp** *ipv6-address* parameter specifies the address of the TFTP server. You must specify this address in hexadecimal using 16-bit values between colons as documented in RFC 2373.

The *source-file-name* parameter specifies the name of the file you want to copy from the TFTP server.

Specify the **running-config** keyword to upload the specified file from the IPv6 TFTP server to the device. The device copies the specified file into the current running configuration but does not overwrite the current configuration.

Specify the **startup-config** keyword to upload the specified file from the IPv6 TFTP server to the device. The the device copies the specified file into the current startup configuration but does not overwrite the current configuration.

Using SNMP to save and load configuration information

You can use a third-party SNMP management application such as HP OpenView to save and load a configuration on a Brocade device. To save and load configuration information using HP OpenView, use the following procedure.

NOTE
The syntax shown in this section assumes that you have installed HP OpenView in the "/usr" directory.

---

1. Configure a read-write community string on the Brocade device, if one is not already configured. To configure a read-write community string, enter the following command from the global CONFIG level of the CLI.

   snmp-server community *string* **ro | rw**

   where *string* is the community string and can be up to 32 characters long.

2. On the Brocade device, enter the following command from the global CONFIG level of the CLI.

   no snmp-server pw-check

   This command disables password checking for SNMP set requests. If a third-party SNMP management application does not add a password to the password field when it sends SNMP set requests to a Brocade device, by default the Brocade device rejects the request.

3. From the command prompt in the UNIX shell, enter the following command.

   /usr/OV/bin/snmpset -c rw-community-string device-ip-addr

   1.3.6.1.4.1.1991.1.1.2.1.5.0

   a tftp-ip-addr 1.3.6.1.4.1.1991.1.1.2.1.8.0 s config-file-name

   1.3.6.1.4.1.1991.1.1.2.1.9.0 integer command-integer

   where

   rw-community-string is a read-write community string configured on the Brocade device.

   fdry-ip-addr is the IP address of the Brocade device.

   tftp-ip-addr is the TFTP server IP address.

   config-file-name is the configuration file name.

   command-integer is one of the following:

   20 – Upload the startup-config file from the flash memory of the Brocade device to the TFTP server.

21 – Download a startup-config file from a TFTP server to the flash memory of the Brocade device.

22 – Upload the running-config from the flash memory of the Brocade device to the TFTP server.

23 – Download a configuration file from a TFTP server into the running-config of the Brocade device.

NOTE
Option **23** adds configuration information to the running-config on the device, and does not replace commands. If you want to replace configuration information in the device, use "no" forms of the configuration commands to remove the configuration information, then use configuration commands to create the configuration information you want. Follow the guidelines in "Dynamic configuration loading" on page 85.

Erasing image and configuration files

To erase software images or configuration files, use the commands described below. These commands are valid at the Privileged EXEC level of the CLI:

- **erase flash primary** erases the image stored in primary flash of the system.
- **erase flash secondary** erases the image stored in secondary flash of the system.
- **erase startup-config** erases the configuration stored in the startup configuration file; however, the running configuration remains intact until system reboot.

System reload scheduling

In addition to reloading the system manually, you can configure the Brocade device to reload itself at a specific time or after a specific amount of time has passed.

NOTE
The scheduled reload feature requires the system clock. Refer to "Network Time Protocol Version 4 (NTPv4)" on page 20.

Reloading at a specific time

To schedule a system reload for a specific time, use the **reload at** command. For example, to schedule a system reload from the primary flash module for 6:00:00 AM, April 1, 2003, enter the following command at the global CONFIG level of the CLI.

```
Brocade#reload at 06:00:00 04-01-03
```

Syntax: **reload at** *hh:mm:ss mm-dd-yy* [**primary** | **secondary**]

hh:mm:ss is the hours, minutes, and seconds.

mm-dd-yy is the month, day, and year.

primary | **secondary** specifies whether the reload is to occur from the primary code flash module or the secondary code flash module. The default is **primary**.

Reloading after a specific amount of time

To schedule a system reload to occur after a specific amount of time has passed on the system clock, use **reload after** command. For example, to schedule a system reload from the secondary flash one day and 12 hours later, enter the following command at the global CONFIG level of the CLI.

```
Brocade#reload after 01:12:00 secondary
```

Syntax: **reload after** *dd:hh:mm* [**primary** | **secondary**]

dd:hh:mm is the number of days, hours, and minutes.

primary | **secondary** specifies whether the reload is to occur from the primary code flash module or the secondary code flash module.

Displaying the amount of time remaining before a scheduled reload

To display how much time is remaining before a scheduled system reload, enter the following command from any level of the CLI.

```
Brocade#show reload
```

Canceling a scheduled reload

To cancel a scheduled system reload using the CLI, enter the following command at the global CONFIG level of the CLI.

```
Brocade#reload cancel
```

Diagnostic error codes and remedies for TFTP transfers

This section describes the error messages associated with TFTP transfer of configuration files, software images or flash images to or from a Brocade device.

Error code	Message	Explanation and action
1	Flash read preparation failed.	A flash error occurred during the download. Retry the download. If it fails again, contact customer support.
2	Flash read failed.	
3	Flash write preparation failed.	
4	Flash write failed.	
5	TFTP session timeout.	TFTP failed because of a time out. Check IP connectivity and make sure the TFTP server is running.

Error code	Message	Explanation and action
6	TFTP out of buffer space.	The file is larger than the amount of room on the device or TFTP server. If you are copying an image file to flash, first copy the other image to your TFTP server, then delete it from flash. (Use the **erase flash...** CLI command at the Privileged EXEC level to erase the image in the flash.) If you are copying a configuration file to flash, edit the file to remove unnecessary information, then try again.
7	TFTP busy, only one TFTP session can be active.	Another TFTP transfer is active on another CLI session or network management system. Wait, then retry the transfer.
8	File type check failed.	You accidentally attempted to copy the incorrect image code into the system. For example, you might have tried to copy a Chassis image into a Compact device. Retry the transfer using the correct image.
16	TFTP remote - general error.	The TFTP configuration has an error. The specific error message describes the error. Correct the error, then retry the transfer.
17	TFTP remote - no such file.	
18	TFTP remote - access violation.	
19	TFTP remote - disk full.	
20	TFTP remote - illegal operation.	
21	TFTP remote - unknown transfer ID.	
22	TFTP remote - file already exists.	
23	TFTP remote - no such user.	

This section describes the error messages associated with the TFTP transfer of PoE firmware file to a Brocade device.

Message	Explanation and action
Firmware TFTP timeout.	TFTP failed because of a time out. Check IP connectivity and make sure the TFTP server is running.
Firmware is not valid for this platform.	Each PoE firmware file delivered by Brocade is meant to be used on the specific platform only. If the file is used on a platform for which it is not meant, then this error message will display. Download the correct file, then retry the transfer.
Firmware is not valid for the IEEE 802.3at (PoE-Plus) controller type.	Each PoE firmware file delivered by Brocade is meant to be used on the specific platform only. If the file is used on a platform for which it is not meant, then this error message will display. Download the correct file, then retry the transfer.
Firmware is not valid for the IEEE 802.3af PoE controller type.	

Message	Explanation and action
Firmware type cannot be detected from the firmware content.	Each PoE firmware file delivered by Brocade is meant to be used on the specific platform and the specific PoE controller on the specified module. If the file is used for a platform for which it is meant, but the PoE controller is not same then this error message will display. Download the correct file, then retry the transfer.
TFTP File not Valid for PoE Controller Type.	
Firmware tftp remote file access failed.	The TFTP server needs read access on the PoE firmware file. Check the permissions on the file, then try again.

Network connectivity testing

After you install the network cables, you can test network connectivity to other devices by pinging those devices. You also can observe the LEDs related to network connection and perform trace routes.

For more information about observing LEDs, refer to the *Brocade FastIron X Series Chassis Hardware Installation Guide* and the *Brocade FastIron Compact Switch Hardware Installation Guide*.

Pinging an IPv4 address

NOTE
This section describes the *IPv4* **ping** command. For details about *IPv6* **ping**, refer to the *FastIron Ethernet Layer 3 Routing Configuration Guide*.

To verify that a Brocade device can reach another device through the network, enter a command such as the following at any level of the CLI on the Brocade device:

```
Brocade> ping 10.33.4.7
```

Syntax: **ping** *ip addr* | *hostname* [**source** *ip addr*] [**count** *num*] [**timeout** *msec*] [**ttl** *num*] [**size** *byte*] [**quiet**] [**numeric**] [**no-fragment**] [**verify**] [**data** *1-to-4 byte hex*] [**brief** [**max-print-per-sec** *number*]]

NOTE
If the device is a Brocade Layer 2 Switch or Layer 3 Switch, you can use the host name only if you have already enabled the Domain Name Server (DNS) resolver feature on the device from which you are sending the ping. Refer to “IP Configuration” chapter in the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide*.

The required parameter is the IP address or host name of the device.

The **source** *ip addr* specifies an IP address to be used as the origin of the ping packets.

The **count** *num* parameter specifies how many ping packets the device sends. You can specify from 1 – 4294967296. The default is 1.

The **timeout** *msec* parameter specifies how many milliseconds the Brocade device waits for a reply from the pinged device. You can specify a timeout from 1 – 4294967296 milliseconds. The default is 5000 (5 seconds).

The **ttl** *num* parameter specifies the maximum number of hops. You can specify a TTL from 1 – 255. The default is 64.

The **size** *byte* parameter specifies the size of the ICMP data portion of the packet. This is the payload and does not include the header. You can specify from 0 – 10000. The default is 16.

The **no-fragment** parameter turns on the "don't fragment" bit in the IP header of the ping packet. This option is disabled by default.

The **quiet** parameter hides informational messages such as a summary of the ping parameters sent to the device and instead only displays messages indicating the success or failure of the ping. This option is disabled by default.

The **verify** parameter verifies that the data in the echo packet (the reply packet) is the same as the data in the echo request (the ping). By default the device does not verify the data.

The **data** *1 – 4 byte hex* parameter lets you specify a specific data pattern for the payload instead of the default data pattern, "abcd", in the packet data payload. The pattern repeats itself throughout the ICMP message (payload) portion of the packet.

NOTE
For numeric parameter values, the CLI does not check that the value you enter is within the allowed range. Instead, if you do exceed the range for a numeric value, the software rounds the value to the nearest valid value.

The **brief** parameter causes ping test characters to be displayed. The following ping test characters are supported:

- **!** Indicates that a reply was received.
- **.** Indicates that the network server timed out while waiting for a reply.
- **U** Indicates that a destination unreachable error PDU was received.
- **I** Indicates that the user interrupted ping.

NOTE
The number of **!** characters displayed may not correspond to the number of successful replies by the **ping** command. Similarly, the number of **.** characters displayed may not correspond to the number of server timeouts that occurred while waiting for a reply. The "success" or "timeout" results are shown in the display as "Success rate is *XX* percent (*X*/*Y*)".

The optional **max-print-per-sec** *number* parameter specifies the maximum number of target responses the Brocade device can display per second while in brief mode. You can specify from 0 – 2047. The default is 511.

NOTE
If you address the ping to the IP broadcast address and network address, the device lists the first four responses to the ping.

NOTE
On 48GC modules in non-jumbo mode, the maximum size of ping packets is 1486 bytes and the maximum frame size of tagged traffic is no larger than 1581 bytes.

Tracing an IPv4 route

NOTE
This section describes the *IPv4* **traceroute** command. For details about *IPv6* **traceroute**, refer to the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide*.

Use the **traceroute** command to determine the path through which a Brocade device can reach another device. Enter the command at any level of the CLI.

The CLI displays trace route information for each hop as soon as the information is received. Traceroute requests display all responses to a given TTL. In addition, if there are multiple equal-cost routes to the destination, the Brocade device displays up to three responses by default.

```
Brocade> traceroute 10.33.4.7
```

Syntax: **traceroute** *host-ip-addr* [**maxttl** *value*] [**minttl** *value*] [**numeric**] [**timeout** *value*] [**source-ip** *ip-addr*]

Possible and default values are as follows.

minttl – minimum TTL (hops) value: Possible values are 1 – 255. Default value is 1 second.

maxttl – maximum TTL (hops) value: Possible values are 1 – 255. Default value is 30 seconds.

timeout – Possible values are 1 – 120. Default value is 2 seconds.

numeric – Lets you change the display to list the devices by their IP addresses instead of their names.

source-ip *ip-addr* – Specifies an IP address to be used as the origin for the traceroute.

Hitless management on the FSX 800 and FSX 1600

Hitless management is supported on the FSX 800 and FSX 1600 chassis with dual management modules. It is a high-availability feature set that ensures no loss of data traffic during the following events:

- Management module failure or role change
- Software failure
- Addition or removal of modules
- Operating system upgrade

During such events, the standby management module takes over the active role and the system continues to forward traffic seamlessly, as if no failure or topology change has occurred. In software releases that do not support hitless management, events such as these could cause a system reboot, resulting in an impact to data traffic.

The following Hitless management features are supported:

Hitless Switchover – A manually controlled (CLI-driven) switchover of the active and standby management modules without any packet loss to the services and protocols that are supported by Hitless management. A switchover is activated by the CLI command **switch-over-active-role**.

Hitless Failover – An automatic, forced switchover of the active and standby management modules because of a failure or abnormal termination of the active management module. In the event of a failover, the active management module abruptly leaves and the standby management module immediately assumes the active role. Like a switchover, a failover occurs without any packet loss to hitless-supported services and protocols. Unlike a switchover, a failover generally happens without warning.

Hitless Operating System (OS) Upgrade – An operating system upgrade and controlled switchover without any packet loss to the services and protocols that are supported by Hitless management.

The services and protocols supported by Hitless management are listed in Table 15 on page 101.

Hitless failover and hitless switchover are disabled by default. To enable these features, refer to "Enabling hitless failover on the FSX 800 and FSX 1600" on page 105 and "Executing a hitless switchover on the FSX 800 and FSX 1600" on page 106.

Benefits of hitless management

The benefits of Hitless management include the following:

- The standby management module (the module that takes over the active role) and all interface modules in the chassis are not reset
- Existing data traffic flows continue uninterrupted with no traffic loss
- Port link states remain UP for the duration of the hitless management event
- System configurations applied through Console/SNMP/HTTP interfaces remain intact
- Hitless switchover can be used by a system administrator, for example, to perform maintenance on a management module that has been functioning as the active management module. Some advantages of a hitless switchover over a hitless software reload are:
    - A manual switchover is quicker, since the standby module does not have to reboot.
    - Switched traffic through the Ethernet interfaces on the standby management module is not interrupted.

NOTE
All traffic going through Ethernet interfaces (if present) on the management modules will be interrupted during a hitless OS upgrade. This is because both management modules must be reloaded with the new image. This applies to hitless OS upgrade only. It does not apply to hitless switchover or failover, which does not interrupt traffic going through Ethernet interfaces on the standby management module (the module that takes over the active role).

Supported protocols and services for hitless management events

Table 15 lists the services and protocols that are supported by Hitless management. Table 15 also highlights the impact of *Hitless management events* (switchover, failover, and OS upgrade) to the system's major functions.

NOTE
Services and protocols that are not listed in Table 15 may be disrupted, but will resume normal operation once the new active management module is back up and running.

TABLE 15 Hitless-supported services and protocols – FSX 800 and FSX 1600

Traffic type	Supported protocols and services	Impact
Layer 2 switched traffic, including unicast and multicast + System-level + Layer 4	• 802.1p and 802.1Q • 802.3ad – LACP • 802.3af – PoE • 802.3at – PoE+ • DSCP honoring and Diffserv • Dual-mode VLAN • IGMP v1, v2, and v3 snooping • IPv4 ACLs • IPv6 ACLs • Layer 2 switching (VLAN and 802.1Q-in-Q) • MLD v1 and v2 snooping • MRP • Multiple spanning tree (MSTP) • Physical port/link state • PIM SM snooping • Port mirroring and monitoring • Port trunking • Rapid spanning tree (RSTP) • Spanning tree (STP) • ToS-based QoS • Policy Based Routing • Traffic policies • UDLD • VSRP	Layer 2 switched traffic is not impacted during a Hitless management event. All existing switched traffic flows continue uninterrupted. New switched flows are not learned by the FastIron switch during the switchover process and are flooded to the VLAN members in hardware. After the new active management module becomes operational, new switched flows are learned and forwarded accordingly. The Layer 2 control protocol states are not interrupted during the switchover process. Configured ACLs, PBR or GRE & IPv6 to IPv4 Tunnels will operate in a hitless manner.
Layer 3 IPv4 routed traffic	• BGP4 • IPv4 unicast forwarding • OSPFv2 • OSPFv2 with ECMP • Static routes • IPv4 PIM (IPv4 non-stop multicast routing needs to be enabled for IPv4 PIM to be hitless.) • VRRP • VRRP-E • GRE • IPv6 to IPv4 Tunnels	Layer 3 routed traffic for supported protocols is not impacted during a Hitless management event. Other Layer 3 protocols that are not supported will be interrupted during the switchover or failover. If BGP4 graceful restart or OSPF graceful restart is enabled, it will be gracefully restarted and traffic will converge to normalcy after the new active module becomes operational. Refer to “OSPF graceful restart” and “BGP4 graceful restart” sections in the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide.* Configured ACLs, PBR or GRE & IPv6 to IPv4 Tunnels will operate in a hitless manner.

TABLE 15 Hitless-supported services and protocols – FSX 800 and FSX 1600

Traffic type	Supported protocols and services	Impact
Layer 3 IPv6 routed traffic	• BGP4+ • IPv6 unicast forwarding • OSPFv3 • OSPFv3 with ECMP • Static routes • VRRP • VRRP-E	Layer 3 routed traffic for supported protocols is not impacted during a Hitless management event. Traffic will converge to normalcy after the new active module becomes operational. Other Layer 3 protocols that are not supported will be interrupted during the switchover or failover. If BGP4+ graceful restart or OSPF graceful restart / OSPFv3 NSR is enabled, it will be gracefully restarted and traffic will converge to normalcy after the new active module becomes operational. For details about OSPFv3 graceful restart, refer to "OSPF v3 graceful restart" section in the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide.* For details about BGP4 graceful restart, refer to "BGP4+graceful restart" section in the *FastIron Ethernet Switch Layer 3 Routing Configuration Guide.* Configured ACLs will operate in a hitless manner.
Management traffic	N/A	All existing management sessions (SNMP, TELNET, HTTP, HTTPS, FTP, TFTP, SSH etc.), are interrupted during the switchover or failover process. All such sessions are terminated and can be re-established after the new Active Controller takes over.
Security	• 802.1X, including use with dynamic ACLs and VLANs • IPv4 ACLs • IPv6 ACLs • DHCP snooping • Dynamic ARP inspection • EAP with RADIUS • IP source guard • Multi-device port authentication, including use with dynamic ACLs and VLANs	Supported security protocols and services are not impacted during a switchover or failover. **NOTE:** If 802.1X and multi-device port authentication are enabled together on the same port, both will be impacted during a switchover or failover. Hitless support for these features applies to ports with 802.1X *only* or multi-device port authentication *only*. Configured ACLs will operate in a hitless manner, meaning the system will continue to permit and deny traffic during the switchover or failover process.
Other services to Management	• AAA • DHCP • sFlow • SNMP v1, v2, and v3 • SNMP traps • NTPv4 • Traceroute	Supported protocols and services are not impacted during a switchover or failover. DNS lookups will continue after a switchover or failover. This information is not synchronized. Ping traffic will be minimally impacted.

Hitless management configuration notes and feature limitations

The following limitations apply to hitless management support.

- All traffic going through Ethernet interfaces (if present) on the management modules will be interrupted during a hitless OS upgrade. This is because both management modules must be reloaded with the new image. This applies to hitless OS upgrade only. It does not apply to hitless switchover or failover, which does not interrupt traffic going through Ethernet interfaces on the standby management module (the module that takes over the active role).
- Static and dynamic multi-slot trunks will flap during a hitless switchover if any of the trunk port members reside on the management module.
- Layer 3 multicast traffic is not supported by Hitless management.

Hitless reload or switchover requirements and limitations

The section describes the design limitation on devices with the following configuration:

- 0-port management modules
- One or more third generation line cards

For hitless reload or switch-over-active-role to succeed, the following requirements and limitations must be met:

- The standby management module must be up and in an "OK {Enabled}" state.
- A configuration requiring a reload must not be pending.
- A hitless-reload must not have already been issued on the previous active management module.
- POE firmware must not be in progress.
- The SXR running configuration must not be classified as too large (greater than 512KB).
- A TFTP session must not be in progress.
- An image sync session must not be in progress.
- The current active management card cannot have a memory utilization of greater than 90% of available memory.
- A line card hotswap must not be in progress.

If any of these conditions are not met, an appropriate error message is printed to the console and hitless-reload or switch-over will not succeed.

What happens during a Hitless switchover or failover

This section describes the internal events that enable a controlled or forced switchover (failover) to take place in a hitless manner, as well as the events that occur during the switchover.

Separate data and control planes

The FSX 800 and FSX 1600 management modules have separate data and control planes. The *data plane* forwards traffic between the switch fabric modules and all of the Interface modules in the chassis. The *control plane* carries traffic that is destined for the CPU of the active management module. Control plane traffic includes the following:

- Management traffic
- Control protocol traffic
- In some cases, the first packet of a data flow

During a controlled or forced switchover, the data plane is not affected. Traffic in the forwarding plane will continue to run without interruption while the standby management module takes over operation of the system. However, traffic in the control plane will be minimally impacted.

Real-time synchronization between management modules

Hitless management requires that the active and standby management modules are fully synchronized at any given point in time. This is accomplished by *baseline* and *dynamic synchronization* of the modules.

When a standby management module is inserted and becomes operational in the FSX 800 or FSX 1600 chassis, the standby module sends a baseline synchronization request to the active management module. The request prompts the active management module to copy the current state of its CPU to the standby CPU, including:

- Start-up and run-time configuration (CLI)
- Layer 2 protocols – Layer 2 protocols such as STP, RSTP, MRP, and VSRP run concurrently on both the active and standby management modules.
- Hardware Abstraction Layer (HAL) – This includes the prefix-based routing table, next hop information for outgoing interfaces, and tunnel information.
- Layer 3 IP forwarding information – This includes the routing table, IP cache table, and ARP table, as well as static and connected routes.
- If NSR is enabled, OSPFv2 and OSPFv3 information is copied to the standby.

As baseline synchronization is performed, the console of the active management module displays the progress of the synchronization.

```
ACTIVE: Detected Stdby heart-beat
ACTIVE: Standby is ready for baseline synchronization.
ACTIVE: Baseline SYNC is completed. Protocol Sync is in progress.
ACTIVE: State synchronization is complete.
```

The first message indicates that the active management module has detected the standby management module. The second message indicates that the standby module has been hot-inserted and is ready for baseline synchronization. The third message is seen when baseline synchronization is completed, and the fourth message is seen when protocol synchronization is completed.

The console of the standby management module also displays the progress of the synchronization.

```
STBY: Baseline SYNC is completed. Protocol Sync is in progress.
STBY: State synchronization is complete.
```

The first message indicates that baseline synchronization is completed, and the second message indicates that protocol sychronization is completed.

When control protocols are synchronized and protocol synchronization timers expire, the standby management module will be in *hot-standby* mode, meaning the standby module is ready to take over as the active management module. In the event of a switchover, the standby module will pick up where the active module left off, without interrupting data traffic.

After baseline synchronization, any new events that occur on the active CPU will be dynamically synchronized on the standby CPU. Examples of such events include:

- CLI/HTTP/SNMP configurations
- CPU receive packets
- Link events
- Interrupts
- Layer 2 and Layer 3 forwarding table updates
- Dynamic user authentication updates such as 802.1X or multi-device port authentication
- Routing protocols OSPFv2 and OSPFv3 updates if NSR is enabled.

Dynamic events are synchronized in such a way that if the active CPU fails before fully executing an event, the standby CPU (newly active CPU) will execute the event after the failover. Also, if the active CPU aborts the event, the standby CPU will abort the event as well.

NOTE
Since both the standby and active management modules run the same code, a command that brings down the active management module will most likely bring down the standby management module. Because all configuration commands are synchronized from active to standby management module in real time, both management modules will reload at almost the same time. This in turn will cause the system to reset all interface modules (similar to the behavior when the **reboot** command is executed) and will cause packet loss associated with a system reboot.

NOTE
If the new active management module becomes out-of-sync with an interface module, information on the interface module can be overwritten in some cases, which can cause an interruption of traffic forwarding.

How a Hitless switchover or failover impacts system functions

Fora description of the feature’s impact to major system functions, refer to Table 15 on page 101.

Enabling hitless failover on the FSX 800 and FSX 1600

Hitless failover is disabled by default. When disabled, the following limitations are in effect:

- If a failover occurs, the system will reload. The following message will display on the console prior to a reload.

  ```
  STBY:- - - - Active Hitless Failover is disabled. Re-setting the system - -
  ```

- Manual switchover (CLI command **switch-over-active-role**) is not allowed. If this command is entered, the following message will display on the console:

  ```
  Switch-over is not allowed. Reason: hitless-failover not configured.
  ```

NOTE
Hitless OS upgrade is *not* impacted by this option and is supported whether or not hitless failover is enabled.

NOTE
Synchronization between the active management module and standby management module will occur whether or not hitless failover is enabled.

To enable hitless failover, enter the following command at the Global CONFIG level of the CLI:

```
Brocade(config)#hitless-failover enable
```

The command takes effect immediately. Manual switchover is allowed, and in the event of a failover, the standby management module will take over the active role without reloading the system.

Syntax: [no] hitless-failover enable

Use the **no** form of the command to disable hitless failover once it has been enabled.

Executing a hitless switchover on the FSX 800 and FSX 1600

Hitless failover must be enabled before a hitless switchover can be executed. To enable hitless failover, refer to "Enabling hitless failover on the FSX 800 and FSX 1600" on page 105.

To switch over to the standby module (and thus make it the active module), enter the following command.

```
Brocade# switch-over-active-role
```

Once you enter this command, the system will prompt you as follows.

```
Are you sure? (enter 'y' or 'n'): y
Running Config data has been changed. Do you want to continue
the switch-over without saving the running config? (enter 'y' or 'n'): n
Please save the running config and try switch-over again
```

Syntax: switch-over-active role

If this command is entered when hitless failover is disabled, the following message will appear on the console:

```
Switch-over is not allowed. Reason: hitless-failover not configured.
```

A management slot which is in active management preference will always attempt to be active on the next reboot.

To reset the preference, enter the command such as the following:

```
Brocade(config)# set-active-mgmt mgmt0/mgmt1
```

Syntax: set-active-management *management slot numbers*

NOTE
The default active management preference is set to mgmt0 (slot 9).

Hitless OS upgrade on the FSX 800 and FSX 1600

Hitless Operating System (OS) Upgrade enables an operating system upgrade and switchover without any packet loss to the services and protocols that are supported by Hitless management.

What happens during a Hitless OS upgrade

The following steps describe the internal events that occur during a hitless OS upgrade.

1. The standby management module resets and reloads with the new software image in its flash memory.
2. The Ethernet interfaces (if present) on the standby module become operational and start carrying data traffic.
3. The active management module synchronizes the standby management module with all the information required to take over the active role.
4. The Layer 2 and Layer 3 control protocols on the standby management module converge. This process takes approximately 70 seconds.
5. The standby management module takes over the active role.
6. The old active management module resets and reloads with the same software image running on the newly active management module.
7. The FastIron switch is now operating with the new software image. The management module that was initially configured as the standby management module is now the active management module and the management module that was initially configured as the active management module is now the standby.

NOTE
The events described above occur internally and do not create or affect the external network topology.

Hitless OS upgrade considerations

Consider the following when using the hitless OS upgrade feature:

- Hitless OS upgrade allows for upgrading the software in a system between two releases of the OS that support this functionality and have compatible data structures. A hitless O/S downgrade may also be supported if the current and target code releases have compatible data structures. From time to time it may be necessary, when enhancing the software or adding new features, to change or add data structures that may cause some releases to be incompatible. In such cases, an upgrade or downgrade will not be hitless, and the software will use the regular Brocade upgrade process - relying on fast reboot.
- For a description of how this feature impacts major system functions, refer to Table 15 on page 101.
- You must have both active and standby management modules installed to use this feature.
- Hitless OS upgrade is supported in software release FSX 05.0.00 or higher, with boot image FSX 05.0.00 or higher. In general, it is supported with patch upgrades, for example, when upgrading from release 07.0.01a to 07.0.01b. It is not supported during major release upgrades, for example when upgrading from release 07.0.00 to 07.1.00.
- This feature can be used to upgrade an image to a higher or lower compatible version of the software. However, if hitless upgrade to a particular software version is not supported, the software upgrade must be performed through a fast reload of the system.
- Hitless OS upgrade between different types of software images is not supported. For example, hitless OS upgrade is supported when upgrading the Layer 2 image to another Layer 2 image. It is not supported when upgrading the Layer 2 image to Layer 3 image, and so on.
- Hitless OS upgrade should be performed locally, since remote connectivity will be lost during the upgrade. During a reload, HTTP, SSH, Telnet, SNMP, and ping sessions will be dropped.
- The active management module switches from the initial active management module to the standby management module during the hitless upgrade process. Therefore, a connection to the console interface on both management modules is required.
- Upon being reset, any traffic going through the ports on the management module will be interrupted. Once the management module is up and running, it will be able to send and receive packets, even before the hitless upgrade process is complete.
- The running configuration is not allowed to be changed any time during the hitless upgrade process.
- System-max configuration changes require a system reload. System-max configuration changes do not take effect by the hitless upgrade. Even if a system-max parameter is changed and saved in the startup configuration, the FastIron switch will revert to the default system-max value upon a hitless software upgrade. The new system-max value will only take effect after a regular system reload.
- Other commands requiring a software reload, such as CAM mode changes, also do not take effect upon hitless upgrade and require a system reload before being placed in effect.

Hitless OS upgrade configuration steps

The following is a summary of the configuration steps for a hitless OS software upgrade.

1. Copy the software image that supports hitless software upgrade from a TFTP server to the FastIron switch. Refer to "Loading the software onto the switch".
2. Install the software image in flash memory on the active and standby management modules.
3. Enter the **hitless-reload** command on the active management module. The command triggers the events described in the section "What happens during a Hitless OS upgrade" on page 107.

Loading the software onto the switch

Hitless OS upgrade loads from the primary and secondary images on the FSX 800 and FSX 1600 Management modules. If you will be using the **hitless-reload** command to perform the hitless upgrade, you must first copy the software image that supports hitless software upgrade onto the flash memory of the active and standby management modules. For instructions, refer to the release notes.

Performing a hitless upgrade

After loading the software image onto the flash memory of the active and standby management modules, you can begin the process of performing a hitless OS upgrade using the **hitless-reload** command. For example,

```
Brocade#hitless-reload primary
```

Syntax: **hitless-reload primary | secondary**

The **primary** parameter specifies that the management module will be reloaded with the **primary** image.

The **secondary** parameter specifies that the management module will be reloaded with the **secondary** image.

NOTE
The **hitless-reload** command is accepted only when the running configuration and startup configuration files match. If the configuration file has changed, you must first save the file (**write mem**) before executing a hitless reload. Otherwise, the following message will display on the console.

```
Error: Running config and start-up config differs. Please reload the system or save
the configuration before attempting hitless reload.
```

Syslog message for Hitless management events

The following Syslog message is generated as a result of a switchover or hitless OS upgrade.

```
SWITCHOVER COMPLETED - by admin - Mgmt Module in slot slotnum is now Active
```

The following Syslog message is generated as a result of a failover.

```
SWITCHOVER COMPLETED - by active CPU failure - Mgmt Module in slot slotnum is now
Active
```

Displaying diagnostic information

Use the following commands to display diagnostic information for a hitless switchover or failover.

```
Brocade#show ipc
Version 6, Grp 0, Recv: stk-p0: 840918, p1: 0, sum: 840918
Message types have callbacks:
 1:Reliable IPC mesage 2:Reliable IPC atomic 4:fragmentation,jumbo
 20:SYNC dynamic change 22:SYNC download reply 24:SYNC download spec i
 25:SYNC restart download 26:SYNC verification 27:SYNC disable/enable
 29:SYNC mgmt hello 35:IPC Ready Msg 36:IPC Msg for Sync Fra
 38:SYNC reliable
Send message types:
 [1]=815798, [21]=1, [35]=1, [38]=24442,
Recv message types:
 [1]=816446,0, [20]=2,0 [22]=1,0
 [29]=25,0, [38]=24442,0,

Statistics:
 send pkt num : 840242, recv pkt num : 840918
 send msg num : 840242, recv msg num : 840918,
 send frag pkt num : 0, recv frag pkt num : 0,
 pkt buf alloc : 832113,

 Reliable-mail   send  success  receive  time us
 target ID         0      0        0        0
 target MAC        0      0        0        0
There is 0 current jumbo IPC session

Possible errors:
 ***recv msg no callback 2, last msg_type=20, from stack0, e1/9
```

Syntax: show ipc

```
Brocade#show ipc_stat
Total available Hsync channel space = 1048580
Total available Appl channel space = 524292
Total number of application msgs in dyn queue = 0
Total number of hsync msgs in dyn queue = 0
Total number of rel sync msgs in dyn queue = 0
Total number of rx pkt msgs in standby dynamic queue
Total number of rx pkt msgs in active dyn queue = 0
Total number of rx pkts relayed = 0
Total number of rx pkts received = 5686578
Total number of dyn-sync messages received so far = 3
Total number of rel-sync pending complete = 0
Total number of L3 baseline-sync packets = 655
Total number of packet drops in sync = 0
Is image_sync_in_progress? = 0
Total num of rx dyn queue drops = 0
Total num of jumbo corrupts = 0
Total number of messages in IP send queue = 0
```

Syntax: show ipc_stat

Displaying management redundancy information

Enter the following command at any level of the CLI, to view the redundancy parameter settings and statistics.

```
Brocade(config)# show redundancy

=== MP Redundancy Settings ===
Configured Active Slot = 9
Running-Config Sync Period = (upon "write mem")

=== MP Redundancy Statistics ===
Current Active Session:
Active mgmt slot = 9, Standby mgmt slot = 10 (Absent)
Switchover cause = No Switchover
Start Time       = Jan  1 00:00:09
Sxr Sys Hitless Enable Status = 0
Total number of Switchover/Failovers = 0
L3 slib baseline sync status: 0  [complete]
```

Layer 3 hitless route purge

Layer 3 traffic is forwarded seamlessly during a failover, switchover, or OS upgrade when hitless management is enabled.

Some protocols support non-stop routing. On enabling non-stop routing, after switchover the management module quickly re-converge the protocol database. Whereas, some protocols support graceful restart, in which the protocol state is re-established with the help of neighboring devices. Once all the protocols converge the routes which were removed from the network during the convergence period, the routes are deleted from the devices. You can set the route purge timer per VRF instance. Configure the timer to set the duration for which the routes should be preserved after switchover. Once this period elapses, the route purging starts, if by then all other protocols have finished non-stop routing or graceful restart.

When switchover occurs, the route purge timer starts. If non-stop routing or graceful restart is also configured, the route validation and purging starts only when they are complete and the purge timer has elapsed. If for some reason more delay is expected in learning the routes, you can configure a larger period for the purge timer.

Setting the IPv4 hitless purge timer on the defatult VRF

To configure the purge timer, enter the **ip hitless-route-purge-timer** command in global configuration mode.

Example for setting IPv4 hitless purge timer on the default VRF

The following example shows how to set the IPv4 hitless purge timer on the default VRF:

```
Brocade(config)# ip hitless-route-purge-timer 60
```

Setting the IPv4 hitless purge timer on the non-default VRF

1. Enter the VRF configuration mode using the **vrf** command.
2. Configure route distinguisher using the **rd** command.
3. Enter IPv4 address family configuration mode using the **address-family ipv4** command.
4. Configure the router purge timer using the **ip hitless-route-purge-timer** command.

Example for setting the IPv4 hitless purge timer on the non-default VRF

The following example shows how to set the IPv4 purge timer on the non-default VRF:

```
Brocade(config)# vrf blue
Brocade(config-vrf-blue)# rd 10:10
Brocade(config-vrf-blue)# address-family ipv4
Brocade(config-vrf-blue-ipv4)# ip hitless-route-purge-timer 60
```

Setting the IPv6 hitless purge timer on the defatult VRF

To configure the purge timer, enter the **ipv6 hitless-route-purge-timer** command in global configuration mode.

Example for setting the IPv6 hitless purge timer on the defatult VRF

The following example shows how to set the IPv6 hitless purge timer on the default VRF:

```
Brocade(config)# ipv6 hitless-route-purge-timer 60
```

Setting the IPv4 hitless purge timer on the non-default VRF

Before you begin: Enable IPv6 unicast routing using the ipv6 unicast-routing command in global configuration mode.

1. Enter the VRF configuration mode using the **vrf** command.
2. Configure route distinguisher using the **rd** command.
3. Enter the IPv6 address family configuration mode using the **address-family ipv6** command.
4. Configure the router purge timer using the **ipv6 hitless-route-purge-timer** command.

Example for setting the IPv6 hitless purge timer on the non-default VRF

The following example shows how to set the IPv6 purge timer on the non-default VRF:

```
Brocade(config)# vrf blue
Brocade(config-vrf-blue)# rd 10:10
Brocade(config-vrf-blue)# address-family ipv6
Brocade(config-vrf-blue-ipv4)# ipv6 hitless-route-purge-timer 60
```

Commands

The following commands supports the features described in this chapter:

- ip hitless-route-purge-timer
- ipv6 hitless-route-purge-timer

ip hitless-route-purge-timer

Configures the maximum time before stale routes are purged from the routing information base (RIB) after a switchover, failover, or OS upgrade. The **no** form of this command sets the purge timer time to its default value.

Syntax

ip hitless-route-purge-timer *seconds*

no ip hitless-route-purge-timer *seconds*

Command Default

By default, the hitless purge timer is set.

Parameters

seconds	Maximum time, in seconds, before stale routes are purged. The valid range is from 2 to 600. The default is 45 seconds.

Command Modes

Global configuration

IPv4 address family configuration

Usage Guidelines

Under normal circumstances, you may not need to change the value of the route purge timer. If you anticipate delay in learning the routes after switchover, you can configure a larger value for the route purge timer.

Examples

The following example shows how to set the IPv4 hitless purge timer on the default VRF:

```
Brocade(config)# ip hitless-route-purge-timer 500
```

The following example shows how to set the IPv4 purge timer on the non-default VRF:

```
Brocade(config)# vrf blue
Brocade(config-vrf-blue)# rd 10:10
Brocade(config-vrf-blue)# address-family ipv4
Brocade(config-vrf-blue-ipv4)# ip hitless-route-purge-timer 120
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands

ipv6 hitless-route-purge-timer

ipv6 hitless-route-purge-timer

Configures the maximum time before stale routes are purged from the routing information base (RIB) after a switchover, failover, or OS upgrade. The **no** form of this command sets the purge timer time to its default value.

Syntax

ipv6 hitless-route-purge-timer *seconds*

no ipv6 hitless-route-purge-timer *seconds*

Command Default

By default, the hitless purge timer is set.

Parameters

seconds	Maximum time, in seconds, before stale routes are purged. The valid range is from 2 to 600. The default is 45 seconds.

Command Modes

Global configuration

IPv6 address family configuration

Usage Guidelines

Under normal circumstances, you may not need to change the value of the route purge timer. If you anticipate delay in learning the routes after switchover, you can configure a larger value for the route purge timer. IPv6 unicast routing must be enabled using the **ipv6 unicast-routing** command before configuring the purge timer.

Examples

The following example shows how to set IPv6 hitless purge timer on default VRF:

```
Brocade(config)# ipv6 hitless-route-purge-timer 500
```

The following example shows how to set IPv6 purge timer on a non-default VRF:

```
Brocade(config)# vrf blue
Brocade(config-vrf-blue)# rd 10:10
Brocade(config-vrf-blue)# address-family ipv6
Brocade(config-vrf-blue-ipv4)# ipv6 hitless-route-purge-timer 120
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands

ip hitless-route-purge-timer

Chapter

Software-based Licensing 4

Table 16 lists the individual Brocade FastIron switches and the software licensing features they support.

TABLE 16 Supported software licensing features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6450
Software-based licensing	Yes	Yes	Yes	Yes
License generation				
License query				
Deleting a license				

This chapter contains the following section:

Software license terminology

This section defines the key terms used in this chapter.

- **Entitlement certificate** – The proof-of-purchase certificate (paper-pack) issued by Brocade when a license is purchased. The certificate contains a unique transaction key that is used in conjunction with the License ID of the Brocade device to generate and download a software license from the Brocade software portal.
- **License file** – The file produced by the Brocade software portal when the license is generated. The file is uploaded to the Brocade device and controls access to a licensed feature or feature set.
- **License ID** (**LID**) – This is a number that uniquely identifies the Brocade device. The LID is used in conjunction with a transaction key to generate and download a software license from the Brocade software portal. The software license is tied to the LID of the Brocade device for which the license was ordered and generated.
- **Licensed feature** – Any hardware or software feature or set of features that require a valid software license in order to operate on the device.

- Transaction key – A unique key, along with the LID, used to generate a software license from the Brocade software portal. The transaction key is issued by Brocade when a license is purchased. The transaction key is delivered according to the method specified when the order is placed:
  - Paper-pack – The transaction key is recorded on an entitlement certificate, which is mailed to the customer.
  - Electronic – The transaction key is contained in an e-mail message, which is sent instantly to the customer after the order is placed. The customer receives the e-mail message within a few minutes after the order is placed, though the timing will vary depending on the network, Internet connection, and so on.

  If a delivery method is not specified at the time of the order, the key will be delivered by the way of paper-pack.

Software-based licensing overview

Prior to the introduction of software-based licensing, Brocade supported hardware-based licensing, where an EEPROM was used to upgrade a Layer 2 or base Layer 3 switch to a premium or advanced Layer 3 switch. With the introduction of software-based licensing, one or more valid software licenses are required to run such licensed features on the device.

Software-based licensing is designed to work together with hardware-based licensing. The first release of software-based licensing employs a combination of hardware-based and software-based licensing. A Brocade device can use hardware-based licensing, software-based licensing, or both. Future releases that support software-based licensing will use software-based licensing only, eliminating the need for a customer- or factory-installed EEPROM on the management module or switch backplane.

Software-based licensing provides increased scalability and rapid deployment of hardware and software features on the supported Brocade family of switches. For example, for premium upgrades, it is no longer necessary to physically open the chassis and install an EEPROM to upgrade the system. Instead, the Web is used to generate, download, and install a software license that enables premium features on the device.

NOTE
An ICX device only supports software-based licensing. Hardware-based licensing is not supported on ICX devices.

How software-based licensing works

A permanent license can be ordered pre-installed in a Brocade device when first shipped from the factory, or later ordered and installed by the customer. In either case, additional licenses can be ordered as needed.

When a license is ordered separately (not pre-installed), an entitlement certificate or e-mail message, along with a transaction key, are issued to the customer by Brocade as proof of purchase. The transaction key and LID of the Brocade device are used to generate a license key from the Brocade software licensing portal. The license key is contained within a license file, which is downloaded to the customer's PC, where the file can then be transferred to a TFTP or SCP server, and then uploaded to the Brocade device.

Once a license is installed on the Brocade device, it has the following effects:

- For FCX and ICX devices, the license unlocks the licensed feature and it becomes available immediately. There is no need to reload the software.
- For FSX devices, the license unlocks the ability to upload the software image (for example, Premium Layer 3) onto the device. Once the software image is installed and the software is reloaded, the license unlocks the licensed feature.
- When a trial license expires, the commands and CLI related to the feature are disabled, but the feature itself cannot be disabled until the system reloads.

Seamless transition for legacy devices

In this chapter, the term *legacy device* refers to a Brocade device that was shipped prior to the introduction of software-based licensing, has an EEPROM installed, and is running pre-release 07.1.00 software.

The transition to software-based licensing is seamless for legacy devices. When upgraded to a release that supports software-based licensing, these devices will continue to operate as previously configured.

NOTE
There are special considerations and instructions for legacy FastIron devices in need of replacement (by way of a Return Merchandise Agreement [RMA]). For details, refer to "Special replacement instructions for legacy devices" on page 143.

License types

The following license types are supported on FastIron devices:

- Application-related – Enables premium or advanced features on the device, for example advanced Layer 3 for the FCX, premium and advanced Layer 3 for ICX 6610, premium Layer 3 for ICX 6450, and full Layer 3 for the FSX 800, and FSX 1600.
- Trial license – Also called a *temporary license,* this enables a license-controlled feature to run on the device on a temporary basis. A trial license enables demonstration and evaluation of a licensed feature and can be valid for a period of 45 days. For more information about a trial license, refer to "Using a trial license" on page 139.
- Normal license – Also called a *permanent license,* this enables a license-controlled feature to run on the device indefinitely.

Non-licensed features

Table 17 lists the FastIron software images that do not require a license to run on the device.

TABLE 17 Software image files that do not require a license

Product	Image filename
FSX 800 FSX 1600	SXSxxxxx.bin SXLxxxxx.bin (base Layer 3) SXLSxxxxx.bin
FCX ICX 6610	FCXSxxxxx.bin (Layer 2) FCXRxxxxx.bin (Layer 3)
ICX 6430 ICX 6450	ICX64Sxxxxx.bin (Layer 2) ICX64Rxxxxx.bin (Layer 3)

For a list of features supported with these images, refer to the release notes.

NOTE
In FSX devices for software release 08.0.00a, you can load an image (see table above) that does not require a license. The image contains some basic Layer 3 functionality; static route, RIP v1, RIP v2, VRRP, and VRRP-E.

Licensed features and part numbers

Table 18 lists the supported licensed features, associated image filenames, and related part numbers.

NOTE
There are no changes to the part numbers for products with pre-installed (factory-installed) licenses. These part numbers are listed for reference in the last column of Table 18.

TABLE 18 Licensed features and part numbers

Product	Licensed feature or feature set	Image filename	Part numbers for software license only	Part numbers for hardware with pre-installed software license
FCX	Advance Layer 3: • BGP4 • GRE	N/A1	FCX-ADV-LIC-SW	FCX624-E-ADV FCX-624-I-ADV FCX624S-ADV FCX624S-HPOE-ADV
				FCX624S-F-ADV
				FCX648-E-ADV FCX648-I-ADV FCX648S-ADV FCX648S-HPOE-ADV

TABLE 18 Licensed features and part numbers (Continued)

Product	Licensed feature or feature set	Image filename	Part numbers for software license only	Part numbers for hardware with pre-installed software license
ICX 6610 Software-based licensing is only supported on ICX 6610 devices.	Premium Layer 3: • OSPF v2 • OSPF v3 • PIM-DM • PIM-SM • PIM-SSM • PIM passive • PBR • RIP v1, v2 • RIPng • VRRP • VRRP v3 for IPv6 • VRRP-E • VRRP-E for IPv6 • IPv6 unicast routing	N/A[1]	ICX 6610-PREM-LIC-SW	ICX 6610-24-PE ICX 6610-24-PI ICX 6610-24P-PE ICX 6610-24P-PI ICX 6610-24F-PE ICX 6610-24F-PI ICX 6610-48-PE ICX 6610-48-PI ICX 6610-48P-PE ICX 6610-48P-PI
	Advance Layer 3: • All features in the Premium license (see the cell above - Premium Layer 3). • BGP4	N/A[1]	ICX 6610-ADV-LIC-SW	Sold separately. To purchase the ICX 6610 Advance license, contact your Brocade representative.
	Upgrade from Premium to Advance license This license is used to upgrade from Premium to Advance license. The license can only be installed on all SKUs that have a Premium license installed.	N/A[1]	ICX 6610-ADV-UPG-LIC-SW	Sold separately. To purchase the Upgrade license, contact your Brocade representative.
	ICX 6610 - Ports on Demand license To upgrade the ICX 6610 1 Gbps ports to 10 Gbps port speed, use the ICX6610-10G-LIC-POD license. By default, the ICX 6610 device has eight active 1 Gbps uplink ports. To increase the uplink capacity of four ports from 1 Gbps to 10 Gbps port speed, purchase a single ICX6610-10G-LIC-POD license. To increase the uplink capacity of all eight ports from 1 Gbps to 10 Gbps port speed, purchase a second ICX6610-10G-LIC-POD license.	N/A[1]	ICX 6610-10G-LIC-POD	Sold separately. To purchase the Ports on Demand license, contact your Brocade representative.
ICX 6450 Software-based licensing is only supported on ICX 6450 devices.	Premium Layer 3: • OSPF v2 • RIP v1, v2 • VRRP • VRRP-E	N/A[1]	ICX6450-PREM-LIC	N/A
	ICX 6450- Ports on Demand license The ICX 6450 device has four active uplink or stacking ports on slot 2. By default, ports 1 and 3 are 10 Gbps ports. By default, without a license at bootup, ports 2 and 4 come up in 10 Gbps port speed in an error disabled state. To enable ports 2 and 4 to 10 Gbps port speed, purchase the ICX6450-2X10G-LIC-POD license. The PoD feature is not applicable to ICX 6430 devices because there are no 10 Gbps ports on the device.	N/A[1]	ICX6450-2X10G-LIC-POD	Sold separately. To purchase the Ports on Demand license, contact your Brocade representative.

TABLE 18 Licensed features and part numbers (Continued)

Product	Licensed feature or feature set	Image filename	Part numbers for software license only	Part numbers for hardware with pre-installed software license
FSX 800 and FSX 1600	IPv4 PREM Layer 3 for IPv4 management modules[1]: • 6,000 active host routes • Anycast RP • BGP4 • DVMRP V2 • IGMP V1, V2, and V3 • ICMP redirect messages • IGMP V3 fast leave (for routing) • MSDP • OSPF V2 • PIM-DM • PIM-SM • PIM passive • Policy-based routing • IPv4 over GRE • Multicast over GRE	SXRxxxxx.bin (full Layer 3)	SX-FIL3U-SW	SX-FI-ZMR-XL
	IPv4 PREM Layer 3 for IPv6-ready management modules: Same features as "IPv4 PREM Layer 3 for IPv4 management modules:"	SXRxxxxx.bin (full Layer 3) SXLRxxxxx.bin	SX-FIL3U-6-IPV4-SW	SX-FI-ZMR-XL SX-FI-ZMR-XL-PREM6 SX-FIZMRXL6
	IPv4 and IPv6 PREM Layer 3 for IPv6-ready management modules: Same features as "IPv4 PREM Layer 3 for IPv4 management modules:", plus the following: • IPv6 Layer 3 forwarding • IPv6 over IPv4 tunnels in hardware • IPv6 redistribution • IPv6 static routes • OSPF V3 • RIPng • IPv6 over GRE	SXRxxxxx.bin (full Layer 3) SXLRxxxxx.bin	SX-FIL3U-6-IPV4-SW and SX-FIL3U-6-IPV6-SW	SX-FI-ZMR-XL SX-FI-ZMR-XL-PREM6 SX-FIZMRXL6

1. Basic Layer 3 functionality is available in an FSX router image that does not require a license, see Table 17.

Table 19 lists the supported software packages.

TABLE 19 Software packages

Product	Software package name	License needed?
FCX	BASE_SOFT_PACKAGE	No
	FCX_FULL_ROUTER_SOFT_PACKAGE	No
	FCX_ADV_ROUTER_SOFT_PACKAGE	Yes
ICX 6610	ICX6610_BASE_ROUTER_SOFT_PACKAGE	No
	ICX6610_PREM_ROUTER_SOFT_PACKAGE	Yes
	ICX6610_ADV_ROUTER_SOFT_PACKAGE	Yes
ICX 6450	ICX6450_BASE_ROUTER_SOFT_PACKAGE	No
	ICX6450_PREM_ROUTER_SOFT_PACKAGE	Yes
FSX 800, FSX 1600	BASE_SOFT_PACKAGE	No
	SX_V4_HW_BASE_L3_SOFT_PACKAGE	No[1]
	SX_V4_HW_ROUTER_SOFT_PACKAGE	Yes
	SX_V6_HW_BASE_L3_SOFT_PACKAGE	No[1]
	SX_V6_HW_ROUTER_IPV4_SOFT_PACKAGE	Yes
	SX_V6_HW_ROUTER_IPV6_SOFT_PACKAGE	Yes

1. Basic Layer 3 functionality is available in an FSX router image that does not require a license, see Table 17.

Licensing rules

This section lists the software licensing rules and caveats related to the Brocade devices that support software-based licensing.

General notes about licensing

The following licensing rules apply to all FastIron devices that support software licensing:

- A license is tied to the unique LID of the management module or fixed configuration switch for which the license was ordered. Therefore, a license can be used on one device only. It cannot be used on any other device.
- More than one license can be installed per device. For example, an FSX 800 with IPv6 hardware can have the license SX-FIL3U-6-IPV4 and the license SX-FIL3U-6-IPV6, and both can be in effect.
- Only one normal or trial license at a time can be in effect for a licensed feature.
- More than one trial license can be in effect at the same time, as long as each trial license applies to a unique licensed feature.
- A trial license cannot replace or supersede a normal license.

Licensing rules for FCX and ICX 6610 devices

The following licensing rules apply to FCX and ICX 6610 devices for software-based licensing. To describe the behavior for running software-based licensing in an FCX traditional stack, the FCX-ADV-LIC-SW license is used as an example.

NOTE
For FCX and ICX 6610 devices, the behavior for running software-based licensing with different licenses (Premium, Advance, or Upgrade licenses) is the same. One license allows multiple protocols to run in a stack. All units must have a separate license to run the same licensed feature in a stack. If all units do not have the same license, the Active controller cannot enable the licensed feature on the stack. If a member unit without a license joins a stack, the Active controller must make sure that no protocols are enabled in a stack before putting a member unit into full operational state.

- Each stack unit in an FCX traditional stack must have a separate software license for the same licensed feature. For example, if there are eight units in a traditional stack, eight separate licenses must be purchased to run the licensed features in the stack. If there is any unit in a stack without the FCX-ADV-LIC-SW license, the Active controller cannot run the licensed features on the stack.
- For example, to run BGP on the stack, the **router bgp** command must enabled through the CLI on the Active controller. If the Active controller does not have the FCX-ADV-LIC-SW license, the user cannot configure the **router bgp** command at the CLI level.
- If the Active controller has the FCX-ADV-LIC-SW license, and the **router bgp** command is enabled at the CLI level, the system checks all operational units to verify that each unit has the FCX-ADV-LIC-SW license. Only if all the operational units have the FCX-ADV-LIC-SW license will the licensed feature run in the stack.
- If any unit does not have the FCX-ADV-LIC-SW license, the **router bgp** command is rejected and the licensed feature cannot run in the stack. For example, the following error message is displayed on the console.

  ```
  Brocade(config)#router bgp
  Error! cannot run BGP because unit 2 has no FCX-ADV license
  ```

- If the Active controller is running BGP (and all other licensed features if enabled), and a unit without the FCX-ADV-LIC-SW license joins the stack, the unit is put into a non-operational state. If a user copies the FCX-ADV-LIC-SW license to a non-operational unit, it takes effect immediately and the unit becomes operational. If the operational unit has a higher priority than the current Standby controller, the unit replaces the existing Standby controller and becomes the new Standby controller. This behavior applies to all cases in which a non-operational unit becomes operational.
- If a user disables BGP from the stack, the Active controller puts all non-operational units in the operational state. The Active controller and the Standby controller must have the same non-operational units. When the Standby controller receives the runtime configuration from the Active controller, the Standby controller must update the state of every unit (operational or non-operational state).
- If a user deletes the FCX-ADV-LIC-SW license in any stack unit, the Active controller does not change the unit to the non-operational state regardless of running BGP or not. Even if a new election algorithm change occurs, an operational unit is not changed to a non-operational unit. The stack continues its BGP running state. The stack cannot run BGP again after the user disables BGP, or after a reload.

- If BGP is not enabled on the Active controller, a stack unit is operational whether or not the Active controller or the stack units have the FCX-ADV-LIC-SW license. This implies that in a stack where all units (Active controller, Standby controller, and member units) have the FCX-ADV-LIC-SW license, a stack can be formed whether or not BGP is enabled. However, if there is a license mismatch between any of the units in a stack, a stack can still be formed provided that the **router bgp** command is not enabled on the Active controller.
- The FCX-ADV-LIC-SW license is not considered when selecting a unit to be the Standby controller.

Licensing rules for FSX 800 and FSX 1600 devices

SX 800 and SX 1600 devices with redundant management modules must have two separate licenses to run the same licensed feature on both management modules. The license file in the active management module is never copied to or updated on the standby management module.

Upon bootup, the active management module compares its license with the standby management module. If the license differs, the active management module immediately shuts down the standby management module. To enable the standby management module, you must install a separate license. For example, if the active management module has the license SX-FIL3U-SW, the standby management module must also have this license.

Licensing for Ports on Demand

NOTE
Licensing for Ports on Demand is applicable to ICX 6610 devices and ICX 6450 devices.

You can use the Ports on Demand (PoD) feature to enable 1 Gbps ports to 10 Gbps port speed.

By default, regardless of what SFP+ media optic is used, the ICX 6610 device has eight active 1 Gbps uplink ports. To increase the uplink capacity of four ports from 1 Gbps to 10 Gbps port speed, purchase a single ICX6610-10G-LIC-POD license. To increase the uplink capacity of all eight ports from 1 Gbps to 10 Gbps port speed, purchase a second ICX6610-10G-LIC-POD license.

The ICX 6450 device has four active uplink and stacking ports on slot 2. By default, regardless of what SFP+ media optic is used, ports 1 and 3 are 10 Gbps ports. By default, without a license at bootup, ports 2 and 4 come up in 10 Gbps port speed in an error disabled state. To enable ports 2 and 4 to 10 Gbps port speed, purchase the ICX6450-2X10G-LIC-POD license. For more information about enabling ports 2 and 4 to 10 Gbps port speed, refer to "Configuration considerations when configuring PoD for ICX 6450 devices only" on page 131.

The PoD feature is not applicable to ICX 6430 devices because there are no 10 Gbps ports on the device.

Configuring PoD on an interface

To upgrade the ICX 6610 and ICX 6450 ports from 1 Gbps to 10 Gbps port speed, perform the following steps.

1. Download the PoD license to the device. For more information about copying the license file on ICX devices, refer to "Using TFTP to copy a license file on FCX and ICX devices" on page 137.
2. Insert the 10 Gbps optic transceiver.

3. Enter the **speed-duplex 10g-full** command on a single, multiple, or interface range on ICX 6610 devices only, as shown in the example below.

```
Brocade(config)# interface ethernet 1/3/1
Brocade(config-if-e10000-1/3/1)# speed-duplex 10g-full
```

NOTE
You do not need to enter the **speed-duplex 10g-full** command on ICX 6450 devices when configuring PoD on an interface.

4. Enter the **write memory** command to save the configuration.

Syntax: [no] speed-duplex [10g-full | 1000-full-master]

The **10g-full** parameter allows you to enable the port speed to 10 Gbps speed.

The **1000-full-master** parameter allows you to enable the port speed to 1 Gbps speed.

If you enable the port to 10 Gbps port speed, and then issue the **no speed-duplex** command, the port continues to run at 10 Gbps speed. Upon bootup, the port reverts to 1 Gbps speed. In a stacking environment, if you issue the **no speed-duplex 10g-full** command on a 10 Gbps port interface, or a range of interfaces, the 10 Gbps port interface defaults back to 1 Gbps port speed in five to ten seconds.

Configuring the upper PoD ports in a stack for ICX 6610 devices only

By default, when a single ICX6610-10G-LIC-POD license is downloaded onto the device, all four lower PoD ports in the stack (1/3/1 to 1/3/4) are eligible for an upgrade to 10 Gbps port speed. If you have a single ICX6610-10G-LIC-POD license, and you want to enable the upper four PoD ports (1/3/5 to 1/3/8) to 10 Gbps port speed, instead of the lower four PoD ports (1/3/1 to 1/3/4), complete the following steps. If the lower four PoD ports are already configured for 10 Gbps speed, you must first issue the **no speed-duplex 10g-full** command on the lower four PoD ports before configuring the upper four PoD ports to 10 Gbps port speed.

The procedure below assumes that you have already downloaded the license.

NOTE
If you purchased a second ICX6610-10G-LIC-POD license, you do not need to perform the steps outlined below. The **flexible-10g-ports upper** command is used to configure the upper four PoD ports to 10 Gbps port speed using a single ICX6610-10G-LIC-POD license. The command can only be used on the upper four PoD ports.

1. Enter the **flexible-10g-ports upper** command at the stack unit configuration level.

```
Brocade(config)# stack unit 2
Brocade(config-unit-2)#flexible-10g-ports upper
Brocade(config-unit-2)#exit
```

2. Specify the upper four PoD ports in a group with a single ICX6610-10G-LIC-POD license at the interface configuration level.

```
Brocade(config)#interface ethernet 2/3/5 to 2/3/8
```

3. Enable the 10 Gbps port speed for the upper four PoD ports by entering the following command.

```
Brocade(config-mif-2/3/5-2/3/8)#speed-duplex 10g-full
Port 2/3/5 mode changed from 1G to 10G
Port 2/3/6 mode changed from 1G to 10G
Port 2/3/7 mode changed from 1G to 10G
Port 2/3/8 mode changed from 1G to 10G
Brocade(config-mif-2/3/5-2/3/8)#end
```

Syntax: [no] **flexible-10g-ports upper**

Use the **no** form of the **flexible-10g-ports upper** command when you want to enable the lower four PoD ports, instead of the upper four PoD ports, to 10 Gbps port speed. Before you issue the **no flexible-10g-ports upper** command, you must first issue the **no speed-duplex 10g-full** command on the upper four PoD ports.

To display the configuration for the **flexible-10g-ports upper** command on a stack unit, use the **show stack** *stack-unit* command.

```
Brocade#show stack 1
alone: standalone, D: dynamic config, S: static config
ID   Type          Role    Mac Address    Pri State   Comment
1  S ICX6610-48    standby 748e.f834.32cc 180 remote  Ready
```

Syntax: **show stack** *stack-unit*

Displaying license configuration for PoD ports after a license upgrade

NOTE
The **show pod** command cannot be used to display the configuration for 1 Gbps or 40 Gbps ports on ICX 6610 and ICX 6450 devices.

The **show pod** command is used to display a license configuration for PoD ports in a stack after a license upgrade on ICX 6610 and ICX 6450 devices. To display general license information about the PoD license in a stack unit, use the **show license** command. For more information about the **show license** command, refer to “Viewing the license database” on page 146.

Displaying license configuration for PoD ports for ICX 6610 devices

To display a license configuration for all PoD ports for all units in a stack, enter the following command at the CLI level.

```
Brocade#show pod
Unit-Id: 1
PoD license capacity:  8
PoD license capacity used:   8

PoD-ports    Lic-Available Lic-Used
1/3/1      Yes             Yes
1/3/2      Yes             Yes
1/3/3      Yes             Yes
```

Syntax: **show pod** [**unit** *id*]

The **unit** *id* parameter specifies the unit ID number of the PoD you want to display.

Table 20 describes the information displayed in the output of the **show pod unit** command.

TABLE 20 Output from the **show pod unit** command

Field	Description
Unit-Id	The unit ID number of the PoD.
PoD license capacity	The port capacity of the PoD license that is purchased. For ICX 6610 devices, the port capacity can be four or eight 10 Gbps ports. UPPER is displayed in parentheses when the upper four PoD ports are selected using the **flexible-10g-ports upper** command. LOWER is displayed in parentheses when the four lower PoD ports are selected for an upgrade to 10 Gbps port speed.
PoD license capacity used	The number of PoD ports that are upgraded to 10 Gbps port speed.
PoD-ports	The list of PoD ports in the PoD unit.
Lic-state	The license state of the PoD ports in the PoD unit.The Lic-state can be one of the following: • default - The port is configured to 1 Gbps speed. • acquired - The port is configured to 10 Gbps speed.

Displaying license configuration for PoD ports for ICX 6450 devices

By default at bootup, the license state for ports 2 and 4 are in the acquired state. The following output from the **show pod** command displays port 2 and 4 as acquired. Upon installing the ICX6450-2X10G-LIC-POD license, ports 2 and 4 can be enabled to run in 10 Gbps port speed. The license state for ports 2 and 4 remains in the acquired state.

```
Brocade#sh pod

Unit-Id: 1

PoD-ports     Lic-Available Lic-Used
1/2/2       Yes              Yes
1/2/4       Yes              Yes

Unit-Id: 2

PoD-ports     Lic-Available Lic-Used
2/2/2       Yes              Yes
2/2/4       Yes              Yes
```

If ports 2 and 4 are configured to 1 Gbps port speed, the license state changes to default. The following output from the **show pod** command displays port 2 and 4 in the default state.

```
Brocade#sh pod

Unit-Id: 1

PoD-ports     Lic-Available Lic-Used
1/2/2       No               No
1/2/4       No               No
```

Syntax: show pod [unit *id*]

The **unit** *id* parameter specifies the unit ID number of the PoD you want to display.

For a description of the fields in the **show pod** and **show pod** [**unit** *id*] command outputs, refer to Table 20 on page 128.

Upgrading or downgrading configuration considerations for PoD

Consider the following when upgrading or downgrading PoD ports for ICX 6610 or ICX 6450 devices:

- When a single ICX6610-10G-LIC-POD license is downloaded onto the device, you can upgrade the first four or the last four PoD ports to 10 Gbps port speed. However, if you upgrade the fifth port to 10 Gbps port speed, the following syslog message and error message displays.

```
SYSLOG: <14>Jul 31 00:33:46 10.20.147.22 PoD: No license present for port
1/3/3".
POD:No license present for port 3.
```

  The error message is displayed because the port you are attempting to upgrade to 10 Gbps port speed has exceeded the license capacity that is downloaded onto the device. To upgrade all eight ICX 6610 ports to 10 Gbps port speed, purchase a second ICX6610-10G-LIC-POD license.

- There is no trial license available for the PoD feature for ICX 6610 and ICX 6450 devices.

Configuration considerations for stacking or trunking PoD ports

Consider the following when stacking or trunking PoD ports for ICX 6610 or ICX 6450 devices:

- In an ICX traditional stack, a stack member unit without a PoD license can join a stack even when the active or master stack unit has a PoD license.
- All trunk ports must operate at 1 Gbps or 10 Gbps speed in a stack. You cannot mix and match trunk ports with different port speeds.
- In a trunk formation, if there is no license installed in a stack upon bootup or hot swap of a stack unit, a port is disabled. This does not affect the trunk formation.

Configuration considerations when configuring PoD on an interface

Consider the following when configuring PoD on an interface or a range of interfaces on ICX 6610 or ICX 6450 devices:

- The **speed-duplex 10g-full** command is rejected when there is no license or an insufficient license in the unit. For example, an insufficient license implies that you are attempting to upgrade all eight ports to 10 Gbps port speed, and you have downloaded a single ICX6610-10G-LIC-POD license. The following syslog message and error message is generated.

```
SYSLOG: <14>Jul 31 00:33:46 10.20.147.22 PoD: No license present for port
1/3/3".
POD:No license present for port 2
```

- When the **speed-duplex 10g-full** command is configured for a port or multiple ports on an interface, and there is no license or an insufficient license in the stack upon bootup or hot swap of a stack unit, the port is configured to 10 Gbps port speed. However, the port is in ERROR_DISABLED state until you install the correct license in the stack. A syslog message is generated every 30 seconds stating that the license is not present on the port. The error disable recovery timer checks every 30 seconds to see if the correct license is installed in the stack. Once the correct license is installed, the port is automatically enabled to operate at 10 Gbps port speed on the next cycle of the timer. You can also enable the port manually to operate at 10 Gbps port speed once the correct license is installed.
- In a stack, the **speed-duplex 10g-full** command is rejected on the primary port if all ports in a trunk do not have the correct license installed to upgrade to 10 Gbps port speed. The following error message is displayed.

```
Error:Not enough 10Gig License present for all the ports in trunk
```

- If you delete the license from the stack, the port runs in 10 Gbps mode until the switch is reloaded. If the **speed-duplex 10g-full** command is entered, the following syslog message is displayed.

```
SYSLOG: <14>Jul 31 00:33:46 10.20.147.22 PoD: No license present for port
1/3/3".
```

- The **show interface ethernet** *stack-unit/slotnum/portnum* command can be used to display the configuration for a disabled 10 Gbps interface port, or a range of port interfaces after bootup for both ICX 6610 and ICX 6450 devices. The **show interface ethernet** command displays the port in the ERROR_DISABLED state. The following example output is from an ICX 6450 device.

```
Brocade#show interface ethernet 3/2/2
10GigabitEthernet3/2/2 is ERR-DISABLED (invalid license), line protocol is
down
  Hardware is 10GigabitEthernet, address is 0000.0083.01fa (bia
0000.0083.01fa)
  Configured speed 10Gbit, actual unknown, configured duplex fdx, actual
unknown
  Stacking Port, port state is DISABLED
  BPDU guard is Disabled, ROOT protect is Disabled
  Link Error Dampening is Disabled
  STP configured to ON, priority is level0, mac-learning is enabled
  Flow Control is enabled
  Mirror disabled, Monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
```

Once the correct license is installed, the port displays the state as up as shown in the example output below.

```
Brocade#show interface ethernet 3/2/2
10GigabitEthernet6/2/2 is up, line protocol is up
  Hardware is 10GigabitEthernet, address is 0000.0082.f872 (bia 0000.0082.f872)

  Configured speed 10Gbit, actual 10Gbit, configured duplex fdx, actual fdx
  Stacking Port, port state is FORWARDING
  BPDU guard is Disabled, ROOT protect is Disabled
  Link Error Dampening is Disabled
  STP configured to ON, priority is level0, mac-learning is enabled
  Flow Control is enabled
  Mirror disabled, Monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
```

Configuration considerations when configuring PoD for ICX 6450 devices only

Consider the following when configuring PoD for ICX 6450 devices only:

- By default, without a license at bootup, ports 2 and 4 come up in 10 Gbps port speed in an error disabled state. The **show interface ethernet** command displays the port in the ERROR_DISABLED state at bootup. Refer to section, "Configuration considerations when configuring PoD on an interface" on page 129 for output example from the **show interface ethernet** command.
- From the default state, ports 2 and 4 can be configured to 1 Gbps port speed using the **speed-duplex 1000-full-master** command without a license. You do not need to reboot the switch for the links to come up in 1 Gbps port speed.
- If you download the ICX6450-2X10G-LIC-POD license to the device, insert the correct 10 Gbps optic transceiver, and enter the **speed-duplex 10g-full** command on the interface, you can immediately begin using ports 2 and 4 in 10 Gbps port speed. You do not need to reboot the switch for the links to come up.
- For any of the four uplink ports on slot 2, if you re-configure any port from 1 Gbps to 10 Gbps port speed, you must reload the switch to begin using the ports in 10 Gbps port speed. Until you reload the switch, the ports will remain in an error-disabled state. The following example output displays ethernet port 4 in an error-disabled state.

```
Brocade#show interface ethernet 1/2/4
10GigabitEthernet1/2/4 is ERR-DISABLED (Reload the switch or stack to enable
this port in 10G speed), line protocol is down
  Hardware is 10GigabitEthernet, address is 0000.0082.e39c (bia 0000.0082.e39c)

  Configured speed 10Gbit, actual unknown, configured duplex fdx, actual
unknown
  Member of L2 VLAN ID 1, port is untagged, port state is DISABLED
  BPDU guard is Disabled, ROOT protect is Disabled
  Link Error Dampening is Disabled
  STP configured to ON, priority is level0, mac-learning is enabled
  Flow Control is enabled
  Mirror disabled, Monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  MTU 1500 bytes
```

Software licensing configuration tasks

This section describes the configuration tasks for generating and obtaining a software license, and then installing it on the Brocade device. Perform the tasks in the order listed in Table 21.

TABLE 21 Configuration tasks for software licensing

	Configuration task	Reference...
1	Order the desired license.	For a list of available licenses and associated part numbers, refer to "Licensed features and part numbers" on page 120.
2	When you receive the transaction key, retrieve the LID of the Brocade device. If you received the transaction key by way of paper-pack, record the LID on the entitlement certificate in the space provided.	"Viewing the License ID" on page 144
3	Log in to the Brocade software portal to generate and obtain the license file.	"Obtaining a license" on page 132
4	Upload the license file to the Brocade device.	"Installing a license file" on page 137
5	Verify that the license is installed.	"Using a trial license" on page 139
6	For FSX 800, FSX 1600, FCX, ICX 6610, and ICX 6450 devices, upload the software image to the device.	Refer to the release notes.

Obtaining a license

The procedures in this section show how to generate and obtain a software license.

1. Order a license for the desired licensed feature. Refer to "Licensed features and part numbers" on page 120 for a list of valid part numbers and licensed features.

   NOTE
   To order and obtain a trial license, contact your Brocade representative.

2. When you receive the paper-pack or electronic transaction key, retrieve the LID of your Brocade device by entering the **show version** command on the device. Example command output is shown in "Viewing the License ID" on page 144."

   If you received a paper-pack transaction key, write the LID in the space provided on the entitlement certificate.

   NOTE
   Do not discard the entitlement certificate or e-mail with electronic key. Keep it in a safe place in case it is needed for technical support or product replacement (RMAs).

3. Log in to the Brocade software portal at *http://swportal.brocade.com* and complete the software license request. If you do not have a login ID and password, request access by following the instructions on the screen.

Figure 2 shows the **Software Portal Login** window.

FIGURE 2 **Brocade Software Portal Login window**

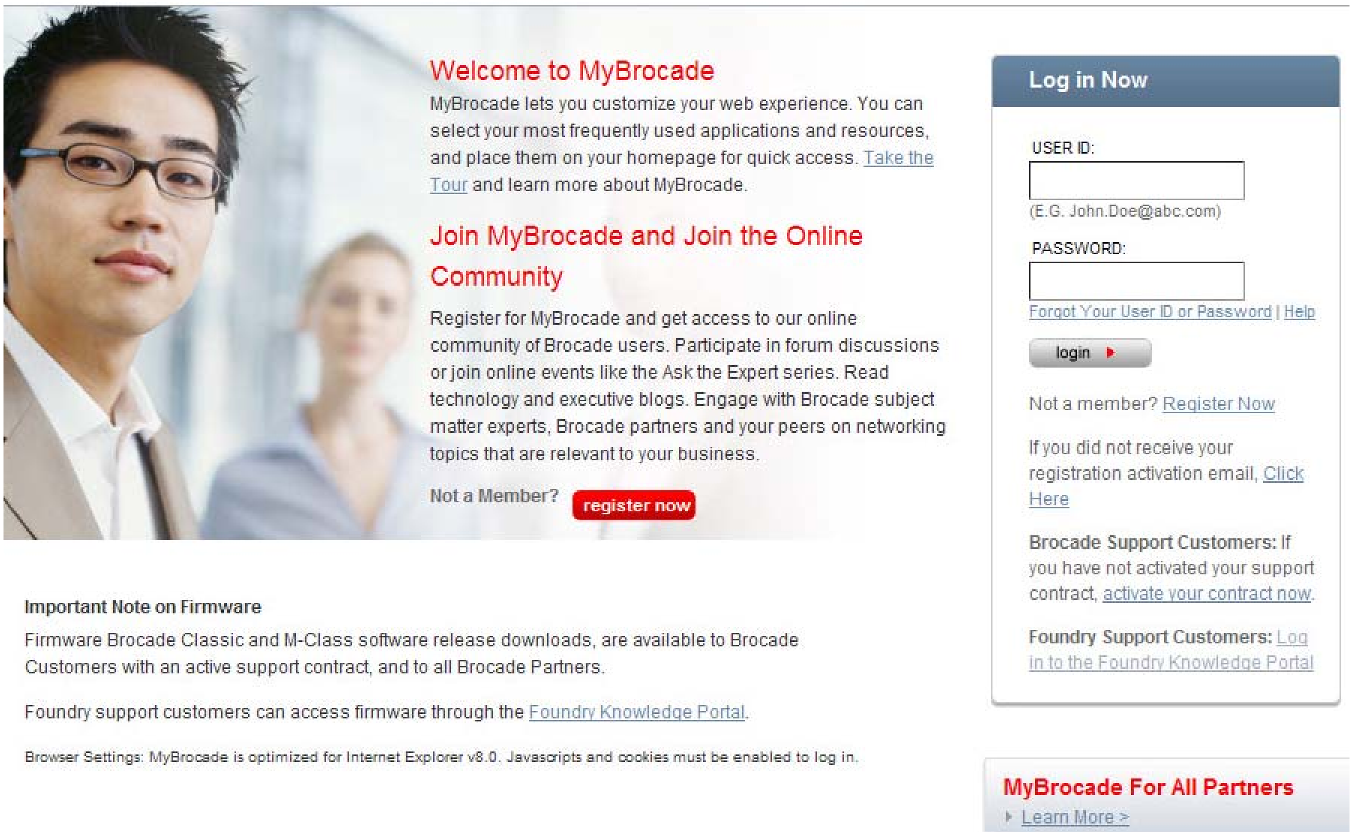

From the **License Management** menu, select **Brocade IP/ADP > License Generation with Transaction key.** The **IP/ADP License Generation** dialog box displays.

FIGURE 3 **License Management Welcome window**

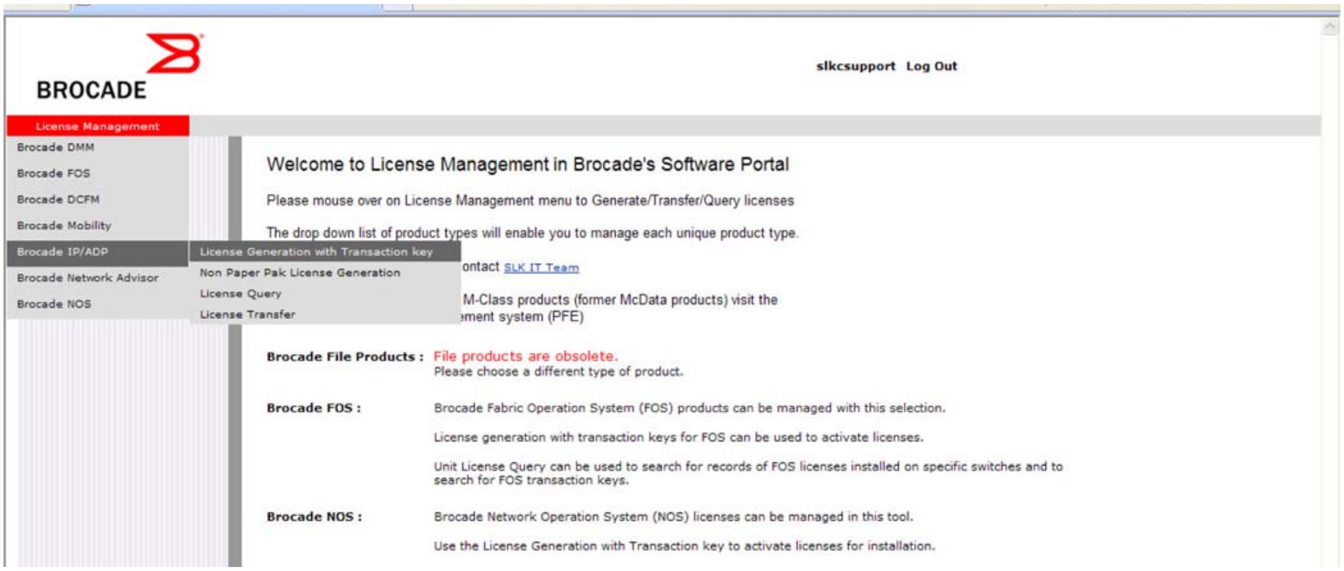

Figure 4 shows the **IP/ADP License Generation** dialog box for generating a license using a transaction key and LID.

FIGURE 4 IP/ADP License Generation window

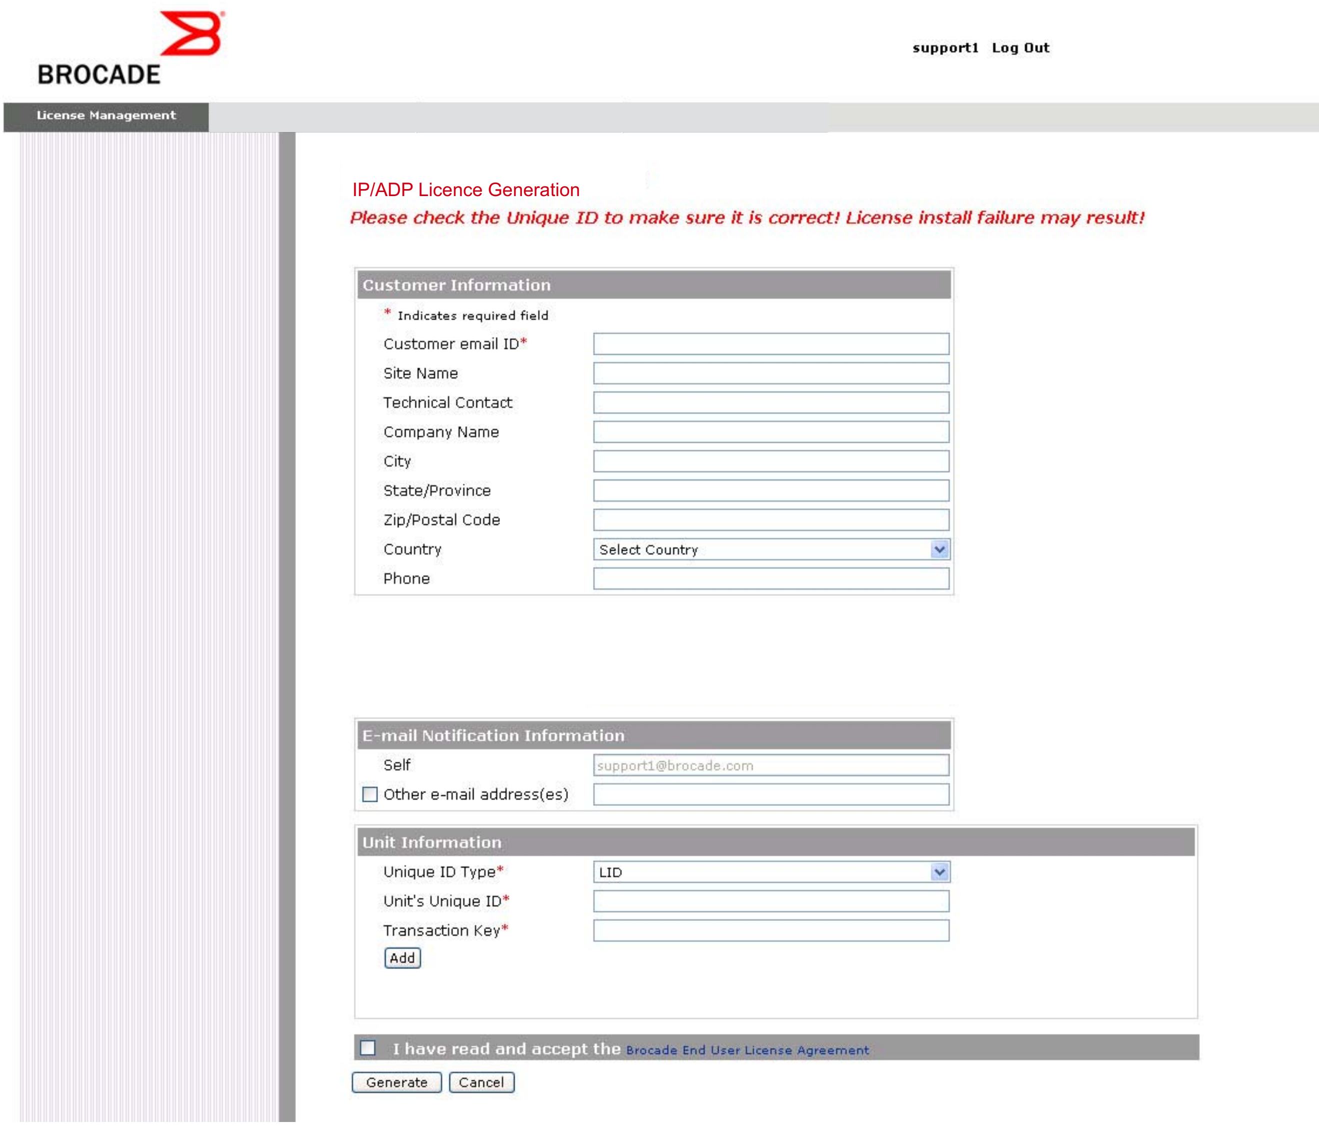

Enter the required information.

- For a description of the field, move the pointer over the field.
- An asterisk next to a field indicates that the information is required.

NOTE
You can generate more than one license at a time. For each license request, enter the **Unit's Unique ID** and **Transaction Key,** and click **Add.**

When you have finished entering the required information, read the Brocade End User License Agreement, and select the **I have read and accept** check box.

Click the **Generate** button to generate the license. Figure 5 shows the results window, which displays an order summary and the results of the license request.

- If the license request is successful, the **Status** field shows “Success” and the **License File** field contains a hyperlink to the generated license file. The license file is automatically sent by e-mail to the specified customer e-mail address.
- If the license request fails, the **Status** field indicates the reason it failed and the action to be taken.

FIGURE 5 **IP/ADP License Generation Results window**

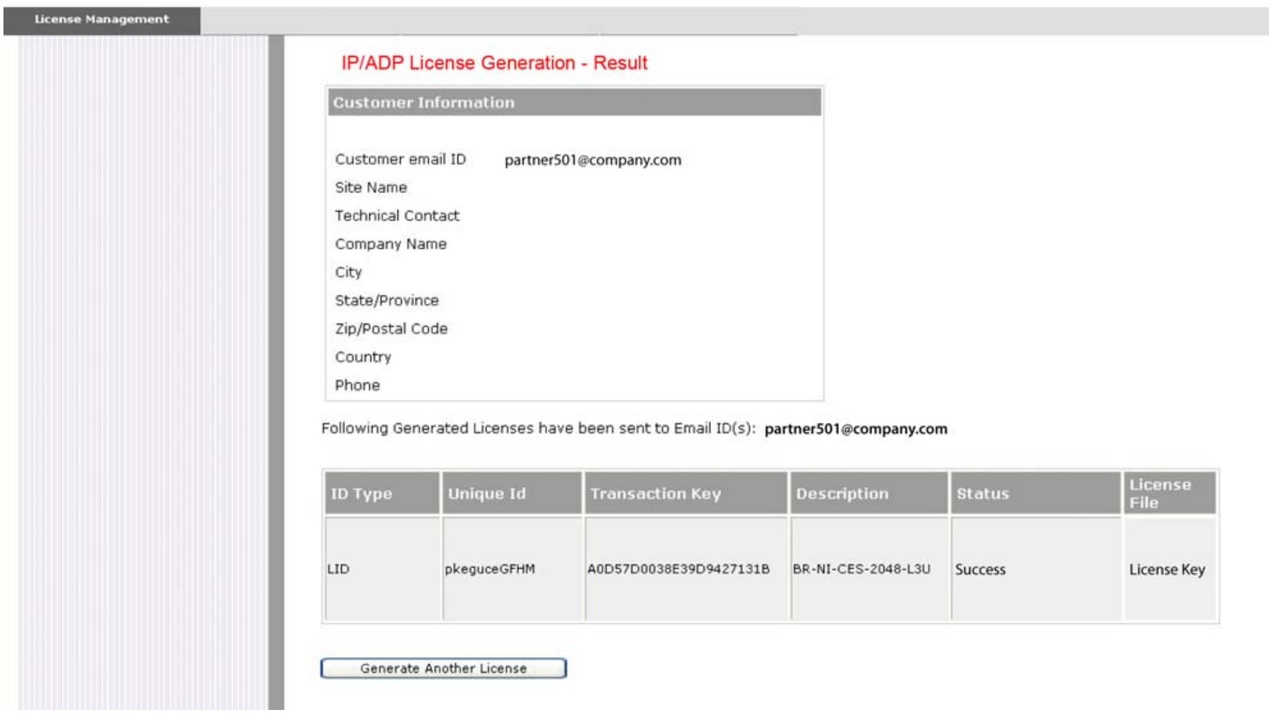

4. Download the license file to your PC by either clicking the hyperlink in the **License File** field or saving the license file from the e-mail attachment.
5. Upload the license file to the Brocade device as instructed in “Installing a license file” on page 137.

Installing a license file

Once you obtain a license file, place it on a TFTP or SCP server to which the Brocade device has access, and then use TFTP or SCP to copy the file to the license database of the Brocade device.

Using TFTP to copy a license file on SX 800 and SX 1600 devices

To copy a license file from a TFTP server to the license database of the Brocade device, enter a command such as the following at the Privileged EXEC level of the CLI:

```
Brocade# copy tftp license 10.1.1.1 lic.xml
```

Syntax: **copy tftp license** [*IP_address* | *ipv6_address*] *license_filename_on_host*

The *IP_address* variable is the address of the IPv4 TFTP server.

The *ipv6_address* variable is the address of the IPv6 TFTP server.

The *license_filename_on_host* variable is the filename of the license file.

If you attempt to download the same license twice on the device, the following error message is displayed on the console.

```
Can't add the license string - 93 (DUPLICATE_LICENSE)
```

Using TFTP to copy a license file on FCX and ICX devices

You can copy a license file from the active unit to all other member units without having to physically disable the stack to install a license for each unit. To copy a license file from the active unit to all other member units in the system, enter a command such as the following at the Privileged EXEC level of the CLI:

```
Brocade# copy tftp license 10.120.54.185 FCX_ADV_LIC_PERP.xml unit 2
```

Syntax: **copy tftp license** [*IP_address* | *ipv6_address*] *license_filename_on_host* **unit** *unit_id*

The *IP_address* variable is the address of the IPv4 TFTP server.

The *ipv6_address* variable is the address of the IPv6 TFTP server.

The *license_filename_on_host* variable is the filename of the license file.

The **unit** *unit_id* parameter specifies a unit for which you want to add a software license file. The *unit_id* variable can be from 1 through 8.

If you attempt to download the same license twice on the device, the following error message is displayed on the console.

```
Can't add the license string - 93 (DUPLICATE_LICENSE)
```

Using Secure Copy to install a license

SSH and Secure Copy (SCP) must be enabled on the Brocade device before the procedures in this section can be performed. For details, refer to the chapter “SSH2 and SCP” in the *FastIron Ethernet Switch Security Configuration Guide*.

The **scp** *license_file_on_host user***@***IP_address***:license** command is supported on SX 800 and SX 1600 devices.

To copy a license file from an SCP-enabled client to the license database of the Brocade device, enter a command such as the following on the SCP-enabled client:

```
c:\scp c:\license\license101 terry@10.1.1.1:license
```

Syntax: **scp** *license_file_on_host user***@***IP_address***:license**

On FCX and ICX devices, to copy a license file from an SCP-enabled client to the license database of a specific unit, enter a command such as the following on the SCP-enabled client:

```
scp license.xml terry@10.20.91.39:license:3
```

In the example above, the license is copied to unit 3.

Syntax: **scp** *license_file_on_host user***@***IP_address***:license:***unit id*

The **unit** *unit_id* parameter specifies a unit for which you want to add a software license file. The *unit_id* variable can be from 1 through 8.

Verifying the license file installation

Use the **show license** command to verify that the license is installed on the device. Details about this command are in "Viewing information about software licenses" on page 144.

Deleting a license file

A license remains in the license database until it is deleted. If you want to delete a license, Brocade recommends that you first disable the licensed feature before deleting the associated license.

NOTE
If a trial license and a normal license for the same licensed feature are added to the system (in either order), the normal license cannot be deleted first. The trial license should be deleted first, and then the normal license can be deleted. If necessary, the same trial license can be reinstalled and used as long as it has not expired.

Deleting a license on SX 800 and SX 1600 devices

To delete a license, enter a command such as the following at the Privileged EXEC level of the CLI:

```
Brocade# license delete 7
```

This command immediately removes the license from the license database. The CLI commands related to the licensed feature will no longer be available from the CLI. The licensed feature will continue to run as configured until the software is reloaded, at which time the feature will be disabled and removed from the system. Syslog and trap messages are generated when the license is deleted.

Syntax: **license delete** *index_number*

The *index_number* variable is a valid license index number. The license index number can be retrieved from the **show license** command output. For more information, refer to "Viewing information about software licenses" on page 144.

NOTE
When removing a license on an FSX device running a release 08.0.00a image, we recommend an immediate reboot.

Deleting a license on FCX and ICX devices

To delete all license files for a specific unit, enter the following command at the Privileged EXEC level of the CLI:

```
Brocade# license delete unit 3 all
```

To delete a specific license file from a unit, enter the following command at the Privileged EXEC level of the CLI:

```
Brocade# license delete unit 3 index 3
```

Syntax: license delete unit *unit_id* [**all** | **index** *license_index*]

The *unit_id* variable specifies the unit ID number.

The **all** option allows you to delete all license files for a specific unit.

The **index** *license_index* parameter specifies the software license file, and is generated by the member unit. The license index number is the license file you want to delete from a unit. The license index number is not unique across stack units, and the user must specify both the unit number and the index number to delete a license from a specific unit. For example, the FCX-ADV-LIC-SW license is installed on both stack unit 3, index 1, and stack unit 5, index 1. Because the index numbers are the same, the user must specify both the unit number and the index number to delete a license from a specific unit.

Using a trial license

NOTE
A trial license must be ordered and installed by a Brocade representative only.

A trial license enables demonstration and evaluation of a licensed feature. The trial license is valid for a period of up to 45 days, and is renewable for an additional 45 days on the second time. A licensed feature operating under a trial license has the same functionality (CLI and **show** commands) as does a licensed feature operating under a normal license. You cannot install more than one trial license of the same license ID on the device. Only one trial license of the same license ID is operational.

What happens when a trial license expires

A trial license expires when it exceeds the specified expiration time or date. The countdown starts when the trial license is generated. When the license expires, the CLI commands related to the licensed feature will no longer be available from the CLI. The licensed feature will continue to run as configured until the system is reloaded, at which time the feature will be disabled and removed from the system.

NOTE
Trial licenses are not cumulative. The new license replaces the current license. To extend the license, you must contact your Brocade representative.

Console, syslog, and trap messages for trial license expiration

Three days prior to the date that a trial license is set to expire, the following warning message will appear daily on the console. On the day that the license will expire, the warning message will appear every two hours. Syslog and trap messages will also be generated.

```
SYSLOG: <12>Jan  1 00:00:00 624-top License: Package FSX624-L3U-IPV6 with LID
NFLLJMI expires in 3 days
```

When the license has expired, the following message will appear on the console. Syslog and trap messages will also be generated.

```
SYSLOG: <13>Jan  1 00:00:00 624-top License: Package FSX624-L3U-IPV6 with LID
NFLLJMI has expired
```

Renewing or extending a trial license

A trial license can be extended once by another trial license of the same type, or by a normal license of the same type. To avoid any interruptions to the network, obtain and install the second trial license before the first license expires. When extended by another trial license, the duration is not cumulative. The countdown starts when the trial license is generated.

To extend the license, you must contact your Brocade representative.

Viewing software license information from the Brocade software portal

This section describes other software licensing tasks supported from the Brocade software portal. You can use the **License Query** option to view software license information for a particular unit, transaction key, or both. You can export the report to Excel for sharing or archiving purposes.

Depending on the status of the license (for example, whether or not the license was generated), the report will include the following Information:

- Hardware part number, serial number, and description
- Software part number, serial number, and description
- Date the license was installed

- Transaction key
- LID
- Feature name
- Product line

From the **License Management** menu, select **Brocade IP/ADP > License Query.**

The **License Query** window displays. (Refer to Figure 6).

FIGURE 6 License Query window

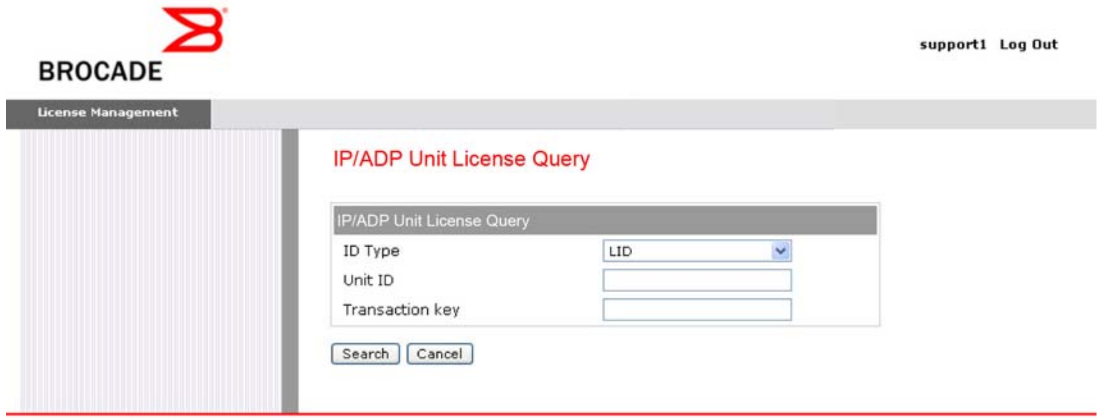

- To view software license information for a particular unit, enter the LID in the **Unit ID** field and click **Search**.
- To view software license information for a particular transaction key, enter the unique number in the **Transaction key** field and click **Search**.

Figure 7 shows an example of the license query results.

FIGURE 7 License Query Results window

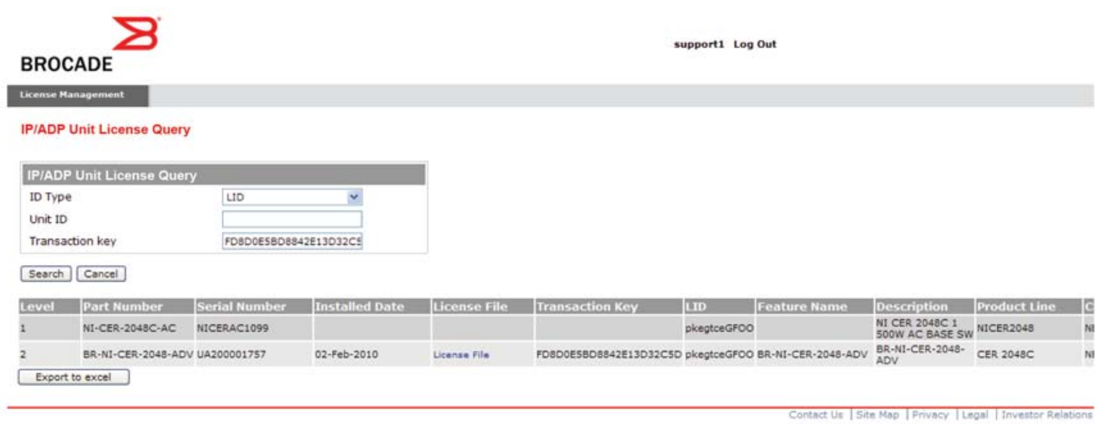

In this example, the line items for Level 1 display hardware-related information and the line items for Level 2 display software-related information. If the query was performed before the transaction key was generated, the first row (Level 1) would not appear as part of the search results. Similarly, if the query was performed before the license was generated, some of the information in the second row would not be displayed.

Transferring a license

A license can be transferred between Brocade devices if both the following conditions are true:

- The device is under an active support contract.
- The license is being transferred between two similar models (for example, from a 24-port model to another 24-port model or from a 48-port model to another 48-port model).

NOTE
Transferring a license is only available internally for TAC, and externally for designated partners with specific accounts in the Software Portal. Contact your Brocade representative for more information.

Special replacement instructions for legacy devices

A legacy device refers to a Brocade device that was shipped prior to the introduction of software-based licensing, has an EEPROM installed, and is running pre-release 07.1.00 software.

For Brocade legacy devices in need of replacement (by way of a Return Merchandise Agreement [RMA]), the following actions must be taken:

- If the replacement device will be upgraded to a software release that supports software-based licensing, registration of the replacement device is required after the software is upgraded.
- If the replacement device will be using a software release that does *not* support software-based licensing, follow these instructions:

  NOTE
  This procedure is not supported on FCX and ICX platforms.

  1. Prior to shipping the device in need of replacement back to the factory, remove the EEPROM from the device. To remove the EEPROM, follow the instructions in the appropriate hardware installation guide or in the instructions that shipped with the EEPROM.
  2. After removing the EEPROM, store it in a safe place.
  3. When the replacement device is received from the factory, install the previously removed EEPROM in the device. To do so, follow the instructions that shipped with the EEPROM.

Syslog messages and trap information

Table 22 lists the syslog messages and traps that are supported for software-based licensing.

TABLE 22 Syslog messages

Message level	Message	Explanation
Informational	License: Package *package_name* with LID *LID_number* is added	The license package has been added.
Informational	License: Package *package_name* with LID *LID_number* is removed	The license package has been deleted.
Warning	License: Package *package_name* with LID *LID_number* expires in *number* days	The trial license is about to expire. This message will begin to display three days before the expiration date, and every two hours on the last day that the license will expire.
Notification	License: Package *package_name* with LID *LID_number* has expired	The trial license has expired.

Viewing information about software licenses

This section describes the **show** commands associated with software-based licensing. These commands are issued on the Brocade device, at any level of the CLI.

NOTE
You can also view information about software licenses from the Brocade software portal. Refer to "Viewing software license information from the Brocade software portal" on page 140.

Viewing the License ID

Brocade devices that ship during and after the release of software-based licensing have the License ID (LID) imprinted on the label affixed to the device. You also can use the **show version** CLI command to view the LID on these devices, and on devices that shipped before the release of software-based licensing.

Use the **show version** command to display the serial number, software and hardware license package name, and LID of all units in the device. The following example is sample output from an FCX unit with the package FCX_ADV_ROUTER_SOFT_PACKAGE installed on units 1, 3, and 5.

```
Brocade#show version
Copyright (c) 1996-2011 Brocade Communications Systems, Inc.
    UNIT 5: compiled on Jun 24 2011 at 14:20:38 labeled as FCXR07203a
                (6674957 bytes) from Primary FCXR07203a.bin
UNIT 1: compiled on Jun 24 2011 at 14:20:38 labeled as FCXR07203a
                (6674957 bytes) from Primary FCXR07203a.bin
UNIT 2: compiled on Jun 24 2011 at 14:20:38 labeled as FCXR07203a
                (6674957 bytes) from Primary FCXR07203a.bin
UNIT 3: compiled on Jun 24 2011 at 14:20:38 labeled as FCXR07203a
                (6674957 bytes) from Primary FCXR07203a.bin
UNIT 4: compiled on Jun 24 2011 at 14:20:38 labeled as FCXR07203a
                (6674957 bytes) from Primary FCXR07203a.bin
Boot-Monitor Image size = 369286, Version:07.0.01T7f5 (grz07001)
```

```
  HW: Stackable FCX648S
==========================================================================
UNIT 1: SL 1: FCX-48GS POE 48-port Management Module
         Serial  #: BCYxxxxxxxxx
         License: FCX_ADV_ROUTER_SOFT_PACKAGE   (LID: deaHHKIgFro)
         P-ENGINE  0: type DB90, rev 01
         P-ENGINE  1: type DB90, rev 01
         PROM-TYPE: FCX-ADV-U
==========================================================================
UNIT 1: SL 2: FCX-2XGC 2-port 16G Module (2-CX4)
==========================================================
===============
UNIT 1: SL 3: FCX-2XG 2-port 10G Module (2-XFP)
==========================================================================
UNIT 2: SL 1: FCX-48GS POE 48-port Management Module
         Serial  #: upgrade7072
         License: FCX_FULL_ROUTER_SOFT_PACKAGE   (LID: ZU0W478MFMH)
         P-ENGINE  0: type DB90, rev 01
         P-ENGINE  1: type DB90, rev 01
==========================================================================
UNIT 2: SL 2: FCX-2XGC 2-port 16G Module (2-CX4)
==========================================================================
==========================================================================
UNIT 3: SL 1: FCX-48GS POE 48-port Management Module
         Serial  #: BCYxxxxxxxxx
         License: FCX_ADV_ROUTER_SOFT_PACKAGE   (LID: deaHHKIgFrN)
         P-ENGINE  0: type DB90, rev 01
         P-ENGINE  1: type DB90, rev 01
==========================================================================
UNIT 3: SL 2: FCX-2XGC 2-port 16G Module (2-CX4)
==========================================================================
==========================================================================
UNIT 4: SL 1: FCX-24GS 24-port Management Module
         Serial  #: BCVxxxxxxxxx
         License: FCX_FULL_ROUTER_SOFT_PACKAGE   (LID: dexHHIIgFFd)
         P-ENGINE  0: type DB90, rev 01
==========================================================================
UNIT 4: SL 2: FCX-2XGC 2-port 16G Module (2-CX4)
==========================================================================
==========================================================================
UNIT 5: SL 1: FCX-48GS 48-port Management Module
         Serial  #: UPGRADE7072
         License: FCX_ADV_ROUTER_SOFT_PACKAGE   (LID: writcfgMFMH)
         P-ENGINE  0: type DB90, rev 01
         P-ENGINE  1: type DB90, rev 01
==========================================================================
```

Syntax: **show version**

In the **show license** command output, only unit 3 and unit 5 are shown with the software license, FCX-ADV-LIC-SW. Unit 1 is not displayed in the **show license** command output because it has a hardware license installed on the device as indicated by the PROM-TYPE: FCX-ADV-U. For more information about the **show license** command, refer to "Viewing the license database" on page 146.

```
Brocade#show license
Index      License Name               Lid          License Type    Status     License Period  License 
Capacity
Stack unit 3:
1         FCX-ADV-LIC-SW           deaHHKIgFrN  Normal         Active    Unlimited                1
Stack unit 4:
1         FCX-ADV-LIC-SW           dexHHIIgFFd  Normal         Active    Unlimited                1
Stack unit 5:
1         FCX-ADV-LIC-SW           writcfgMFMH  Normal         Active    Unlimited                1
```

Syntax: **show license**

Viewing the license database

NOTE
The **show license** command can be used to display software license information for the SX 800 and SX 1600, FCX, ICX 6610, and ICX 6450 devices.

To display general information about all software licenses for all units in a device, use the **show license** command. The **show license** command only displays software license information for a unit, not hardware license information, as shown in the following example.

```
Brocade#show license
Index      License Name Lid           License Type    Status License Period  License Capacity
Stack unit 3:
1 FCX-ADV-LIC-SW deaHHKIgFrN  Normal          Active     Unlimited     1
Stack unit 5:
1 FCX-ADV-LIC-SW writcfgMFMH  Normal          Active     Unlimited     1
```

To display software license information on an ICX 6610 device (for example, the ICX 6610 premium and advance licenses) enter the following command.

```
Brocade#show license
Index      License Name             Lid          License Type    Status     License Period  License Capacity
Stack unit 1:
1          ICX6610-PREM-LIC-SW      FJdnjFJFGiF  Normal          Active     Unlimited                 1
2          ICX6610-10G-LIC-POD      FJdnjFJFGiF  Normal          Active     Unlimited                 8
Stack unit 2:
1          ICX6610-ADV-LIC-SW       FJdnmFJFGiF  Normal          Active     Unlimited                 1
2          ICX6610-10G-LIC-POD      FJdnjFJFGiF  Normal          Invalid    Unlimited                 8
Stack unit 3:
1          ICX6610-ADV-LIC-SW       FJdnlFJFGiF  Normal          Active     Unlimited                 1
4          ICX6610-10G-LIC-POD      FJdnlFJFGiF  Normal          Active     Unlimited                 8
```

Syntax: **show license**

To display software license information on an ICX 6430 device, enter the following command. In the example below, the premium and PoD licenses are installed on stack unit 1, and on stack unit 2 only the premium license is installed.

```
Brocade#show license
Index     License Name              Lid          License Type    Status     License Period  License Capacity
Stack unit 1:
1          ICX6450-PREM-LIC-SW       dbtFJIKiFFI  Normal           Active     Unlimited                   2
2          ICX6450-10G-LIC-POD       dbtFJIKiFFI  Normal           Active     Unlimited                   2
Stack unit 2:
1          ICX6450-PREM-LIC-SW       dbtFJIKiFFI  Normal           Active     Unlimited                   2
```

Syntax: show license

NOTE

The output from the **show license** command is the same for an SX 800 or SX 1600 device.

```
Brocade#show license
Index      Package Name         Lid          License Type     Status     License
Period
1         FCX624-EL3U          cpFNJHFFGO   normal           active     unlimited
```

Syntax: show license

To display specific software license information installed on a SX 800 or SX 1600 device, enter the following command.

```
Brocade#show license 1
License information for license <1>:
        +package name:        FCX624-EL3U
        +lid:                 cpFNJHFFGO
        +license type:        normal
        +status:              active
        +license period:      unlimited
```

Syntax: show license *index_number*

The *index_number* variable specifies the specific license file installed on the device. The **unit** *unit_id* parameter is not applicable on a SX 800 or SX 1600 device.

To display software license information for a specific stack unit on an ICX 6610, ICX 6450, or a FCX device, enter the following command. In the output below, the ICX 6610 premium license, and the POD license are installed on unit 3.

```
Brocade#show license unit 3
Index      License Name              Lid          License Type    Status     License Period  License Capacity
Stack unit 3:
1          ICX6610-PREM-LIC-SW       FJdnjFJFGiF  Normal           Active     Unlimited                   1
2          ICX6610-10G-LIC-POD       FJdnjFJFGiF  Normal           Active     Unlimited                   8
```

Syntax: show license [**unit** *unit_id*]

The **unit** *unit_id* parameter specifies the unit ID number. The unit ID number is available only on FCX, ICX 6610, and ICX 6450 devices.

Table 23 describes the information displayed by the **show license unit** *unit_id* command

TABLE 23 Output from the **show license unit** command

Field	Description
Index	The index number specifies the software license file for a specific stack The index number is generated by the member unit.
License Name	The name of license installed for the license index number on the stack unit.
Lid	The license ID. This number is embedded in the Brocade device.
License Type	Indicates whether the license is normal (permanent) or trial (temporary).
Status	Indicates the status of the license: • **Valid** – A license is valid if the LID matches the serial number of the device for which the license was purchased, and the package name is recognized by the system. • **Invalid** – The LID does not match the serial number of the device for which the license was purchased. • **Active** – The license is valid and in effect on the device. • **Not used** – The license is not in effect on the device. • **Expired** – For trial licenses only, this indicates that the trial license has expired.
License Period	If the license type is trial (temporary), this field displays the number of days the license is valid. If the license type is normal (permanent), this field displays Unlimited.
License capacity	The port capacity of the PoD license. For ICX 6610 devices, the PoD license can be a 4 port capacity license, or a 8 port capacity license depending on the number of licenses purchased. For ICX 6450 devices, the field displays license capacity 2 when the PoD license is purchased for two ports (ports two and four). The license capacity field is displayed in the **show license** command outputs for ICX 6610, ICX 6450, and FCX devices only.

To display detailed information about a license for a specific unit, use the **show license unit** *unit_id* [**index** *index_number*] command. The following example shows a sample output.

```
Brocade#show license unit 3 index 1
License information for unit 3 license <1>:
        +license name:        FCX-ADV-LIC-SW
        +lid:                 deaHHKIgFrN
        +license type:        Normal
        +status:              Active
        +license period:      Unlimited
Trial license information:
        +days used:          10
        +hours used:         21
        +days left:          30
        +hours left:         18
```

Syntax: show license unit *unit_id* [**index** *index_number*]

The **index** *license_index* parameter specifies the software license file that you want to display information for. The **index** *index_number* option is available only on FCX, ICX 6610, and ICX 6450 devices.

Table 24 describes the information displayed by the **show license unit** *unit_id* [**index** *index_number*] command.

TABLE 24 Output from the **show license** command

Field	Description
+license name	The name of the license installed on the unit.
+lid	The license ID. This number is embedded in the Brocade device.
+license type	Indicates whether the license is normal (permanent) or trial (temporary).
+status	Indicates the status of the license: • **Valid** – A license is valid if the LID matches the serial number of the device for which the license was purchased, and the package name is recognized by the system. • **Invalid** – The LID does not match the serial number of the device for which the license was purchased. • **Active** – The license is valid and in effect on the device. • **Not used** – The license is not in effect on the device. • **Expired** – For trial licenses only, this indicates that the trial license has expired.
+license period	If the license type is trial (temporary), this field displays the number of days the license is valid. If the license type is normal (permanent), this field displays Unlimited.
+ days used	The number of days the trial license has been in effect.
+ hours used	The number of hours the trial license has been in effect.
+ days left	The number of days left before the trial license expires.
+ hours left	The number of hours left before the trial license expires.

Viewing software packages installed in the device

Use the **show version** command to view the software packages that are currently running in the device. The following example is sample output from an FCX unit with the package FCX_ADV_ROUTER_SOFT_PACKAGE installed on units 1, 3, and 5.

NOTE
The software package name is not the same as the license name.

```
Brocade#show version
Copyright (c) 1996-2011 Brocade Communications Systems, Inc.
    UNIT 5: compiled on Jun 24 2011 at 14:20:38 labeled as FCXR07203a
                (6674957 bytes) from Primary FCXR07203a.bin
UNIT 1: compiled on Jun 24 2011 at 14:20:38 labeled as FCXR07203a
                (6674957 bytes) from Primary FCXR07203a.bin
UNIT 2: compiled on Jun 24 2011 at 14:20:38 labeled as FCXR07203a
                (6674957 bytes) from Primary FCXR07203a.bin
UNIT 3: compiled on Jun 24 2011 at 14:20:38 labeled as FCXR07203a
                (6674957 bytes) from Primary FCXR07203a.bin
UNIT 4: compiled on Jun 24 2011 at 14:20:38 labeled as FCXR07203a
                (6674957 bytes) from Primary FCXR07203a.bin
Boot-Monitor Image size = 369286, Version:07.0.01T7f5 (grz07001)
  HW: Stackable FCX648S
==========================================================================
UNIT 1: SL 1: FCX-48GS POE 48-port Management Module
         Serial  #: BCYxxxxxxxx
         License: FCX_ADV_ROUTER_SOFT_PACKAGE   (LID: deaHHKIgFro)
         P-ENGINE  0: type DB90, rev 01
```

```
         P-ENGINE  1: type DB90, rev 01
         PROM-TYPE: FCX-ADV-U
==========================================================================
UNIT 1: SL 2: FCX-2XGC 2-port 16G Module (2-CX4)
==========================================================
===============
UNIT 1: SL 3: FCX-2XG 2-port 10G Module (2-XFP)
==========================================================================
UNIT 2: SL 1: FCX-48GS POE 48-port Management Module
         Serial  #: upgrade7072
         License: FCX_FULL_ROUTER_SOFT_PACKAGE   (LID: ZU0W478MFMH)
         P-ENGINE  0: type DB90, rev 01
         P-ENGINE  1: type DB90, rev 01
==========================================================================
UNIT 2: SL 2: FCX-2XGC 2-port 16G Module (2-CX4)
==========================================================================
==========================================================================
UNIT 3: SL 1: FCX-48GS POE 48-port Management Module
         Serial  #: BCYxxxxxxxx
         License: FCX_ADV_ROUTER_SOFT_PACKAGE   (LID: deaHHKIgFrN)
         P-ENGINE  0: type DB90, rev 01
         P-ENGINE  1: type DB90, rev 01
==========================================================================
UNIT 3: SL 2: FCX-2XGC 2-port 16G Module (2-CX4)
==========================================================================
==========================================================================
UNIT 4: SL 1: FCX-24GS 24-port Management Module
         Serial  #: BCVxxxxxxxx
         License: FCX_FULL_ROUTER_SOFT_PACKAGE   (LID: dexHHIIgFFd)
         P-ENGINE  0: type DB90, rev 01
==========================================================================
UNIT 4: SL 2: FCX-2XGC 2-port 16G Module (2-CX4)
==========================================================================
==========================================================================
UNIT 5: SL 1: FCX-48GS 48-port Management Module
         Serial  #: UPGRADE7072
         License: FCX_ADV_ROUTER_SOFT_PACKAGE   (LID: writcfgMFMH)
         P-ENGINE  0: type DB90, rev 01
         P-ENGINE  1: type DB90, rev 01
==========================================================================
```

Syntax: **show version**

For a list of supported software packages installed on the device, refer to Table 19 on page 123.

Chapter

IPv6

5

Table 1 lists the individual Brocade FastIron switches and the IPv6 features they support. These features are supported with premium IPv6 devices running the full Layer 3 software image.

TABLE 1 Supported IPv6 features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6450	ICX 6430
Static IPv6 Routing	Yes [a]	Yes	Yes	Yes	No
IPv6 over IPv4 tunnels	Yes	Yes	Yes	No	No
ECMP load sharing	Yes	Yes	Yes	Yes	No

a. Third generation modules.

This chapter contains the following sections:

Static IPv6 route configuration

NOTE
Static IPv6 route configuration is supported only with the IPv6 Layer 3 license on FSX devices and the full Layer 3 image on other devices.

You can configure a static IPv6 route to be redistributed into a routing protocol, but you cannot redistribute routes learned by a routing protocol into the static IPv6 routing table.

NOTE
The maximum IPv6 static routes supported on an ICX 6450 device is 1070.

Before configuring a static IPv6 route, you must enable the forwarding of IPv6 traffic on the Layer 3 switch using the **ipv6 unicast-routing** command and enable IPv6 on at least one interface by configuring an IPv6 address or explicitly enabling IPv6 on that interface. For more information on performing these configuration tasks, refer to “Configuring IPv4 and IPv6 protocol stacks” section in the *FastIron Ethernet Switch Administration Guide*.

Configuring a static IPv6 route

To configure a static IPv6 route for a destination network with the prefix 2001:DB8::0/32, a next-hop gateway with the global address 2001:DB8:0:ee44::1, and an administrative distance of 110, enter the following command.

```
Brocade(config)#ipv6 route 2001:DB8::0/32 2001:DB8:2343:0:ee44::1 distance 110
```

Syntax: **ipv6 route** *dest-ipv6-prefix*/*prefix-length* *next-hop-ipv6-address* [*metric*] [**distance** *number*]

To configure a static IPv6 route for a destination network with the prefix 2001:DB8::0/32 and a next-hop gateway with the link-local address fe80::1 that the Layer 3 switch can access through Ethernet interface 1/3/1, enter the following command.

```
Brocade(config)#ipv6 route 2001:DB8::0/32 ethernet 1/3/1 fe80::1
```

Syntax: **ipv6 route** *dest-ipv6-prefix*/*prefix-length* [**ethernet** *slot*/*port* | **ve** *num*] *next-hop-ipv6-address* [*metric*] [**distance** *number*]

To configure a static IPv6 route for a destination network with the prefix 2001:DB8::0/32 and a next-hop gateway that the Layer 3 switch can access through tunnel 1, enter the following command.

```
Brocade(config)#ipv6 route 2001:DB8::0/32 tunnel 1
```

Syntax: **ipv6 route** *dest-ipv6-prefix*/*prefix-length* *interface port* [*metric*] [**distance** *number*]

Table 2 describes the parameters associated with this command and indicates the status of each parameter.

TABLE 2 Static IPv6 route parameters

Parameter	Configuration details	Status
The IPv6 prefix and prefix length of the route's destination network.	You must specify the *dest-ipv6-prefix* parameter in hexadecimal using 16-bit values between colons as documented in RFC 2373. You must specify the *prefix-length* parameter as a decimal value. A slash mark (/) must follow the *ipv6-prefix* parameter and precede the *prefix-length* parameter.	Mandatory for all static IPv6 routes.
The route's next-hop gateway, which can be one of the following: • The IPv6 address of a next-hop gateway. • A tunnel interface.	You can specify the next-hop gateway as one of the following types of IPv6 addresses: • A global address. • A link-local address. If you specify a global address, you do not need to specify any additional parameters for the next-hop gateway. If you specify a link-local address, you must also specify the interface through which to access the address. You can specify one of the following interfaces: • An Ethernet interface. • A tunnel interface. • A virtual interface (VE). If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a VE or tunnel interface, also specify the VE or tunnel number. You can also specify the next-hop gateway as a tunnel interface. If you specify a tunnel interface, also specify the tunnel number.	Mandatory for all static IPv6 routes.
The route's metric.	You can specify a value from 1 – 16.	Optional for all static IPv6 routes. (The default metric is 1.)
The route's administrative distance.	You must specify the **distance** keyword and any numerical value.	Optional for all static IPv6 routes. (The default administrative distance is 1.)

A metric is a value that the Layer 3 switch uses when comparing this route to other static routes in the IPv6 static route table that have the same destination. The metric applies only to routes that the Layer 3 switch has already placed in the IPv6 static route table.

The administrative distance is a value that the Layer 3 switch uses to compare this route with routes from other route sources that have the same destination. (The Layer 3 switch performs this comparison before placing a route in the IPv6 route table.) This parameter does not apply to routes that are already in the IPv6 route table. In general, a low administrative distance indicates a preferred route. By default, static routes take precedence over routes learned by routing protocols. If you want a dynamic route to be chosen over a static route, you can configure the static route with a higher administrative distance than the dynamic route.

Configuring a static route in a non-default VRF or User VRF

To configure a static IPv6 route for a destination network with the prefix 2001:DB8::0/32, a next-hop gateway with the global address 2001:DB8:0:ee44::1, in the non-default VRF named "blue", enter the following at the general configuration prompt.

```
Brocade(config)# ipv6 route vrf blue 2001:DB8::0/32 2001:DB8:0:ee44::1
```

Syntax: [**no**] **ipv6 route vrf** *vrf-name dest-ipv6-prefix*/*prefix-length next-hop-ipv6-address*

The *dest-ip-addr* is the route's destination. The *dest-mask* is the network mask for the route's destination IPv6 address.

The *vrf-name* is the name of the VRF that contains the next-hop router (gateway) for the route.

The *next-hop-ip-addr* is the IPv6 address of the next-hop router (gateway) for the route.

NOTE
The **vrf** needs to be a valid VRF to be used in this command.

IPv6 over IPv4 tunnels

NOTE
This feature is supported only with the IPv6 Layer 3 license on FSX devices and the full Layer 3 image on other devices.

To enable communication between isolated IPv6 domains using the IPv4 infrastructure, you can manually configure IPv6 over IPv4 tunnels that provide static point-point connectivity.

As shown in Figure 1, these tunnels encapsulate an IPv6 packet within an IPv4 packet.

FIGURE 1 **IPv6 over an IPv4 tunnel**

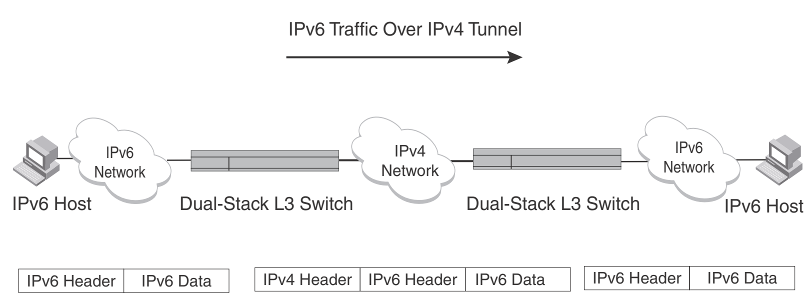

In general, a manually configured tunnel establishes a permanent link between switches in IPv6 domains. A manually configured tunnel has explicitly configured IPv4 addresses for the tunnel source and destination.

This tunneling mechanism requires that the Layer 3 switch at each end of the tunnel run both IPv4 and IPv6 protocol stacks. The Layer 3 switches running both protocol stacks, or dual-stack routers, can interoperate directly with both IPv4 and IPv6 end systems and routers. Refer to "Configuring IPv4 and IPv6 protocol stacks" section in the *FastIron Ethernet Switch Administration Guide.*

IPv6 over IPv4 tunnel configuration notes

- The local tunnel configuration must include both source and destination addresses.
- The remote side of the tunnel must have the opposite source/destination pair.
- A tunnel interface supports static and dynamic IPv6 configuration settings and routing protocols.
- Duplicate Address Detection (DAD) is not currently supported with IPv6 tunnels. Make sure tunnel endpoints do not have duplicate IP addresses.
- Neighbor Discovery (ND) is not supported with IPv6 tunnels.
- If a tunnel source port is a multi-homed IPv4 source, the tunnel will use the first IPv4 address only. For proper tunnel operation, use the **ip address** option.

Configuring a manual IPv6 tunnel

You can use a manually configured tunnel to connect two isolated IPv6 domains. You should deploy this point-to-point tunnelling mechanism if you need a permanent and stable connection.

To configure a manual IPv6 tunnel, enter commands such as the following on a Layer 3 Switch running both IPv4 and IPv6 protocol stacks on each end of the tunnel.

```
Brocade(config)#interface tunnel 1
Brocade(config-tnif-1)#tunnel source ethernet 1/3/1
Brocade(config-tnif-1)#tunnel destination 10.162.100.1
Brocade(config-tnif-1)#tunnel mode ipv6ip
Brocade(config-tnif-1)#ipv6 enable
```

This example creates tunnel interface 1 and assigns a link local IPv6 address with an automatically computed EUI-64 interface ID to it. The IPv4 address assigned to Ethernet interface 1/3/1 is used as the tunnel source, while the IPv4 address 10.168.100.1 is configured as the tunnel destination. The tunnel mode is specified as a manual IPv6 tunnel. Finally, the tunnel is enabled. Note that instead of entering **ipv6 enable**, you could specify an IPv6 address, for example, **ipv6 address 2001:DB8:384d:34::/64 eui-64**, which would also enable the tunnel.

Syntax: [no] **interface tunnel** *number*

For the *number* parameter, specify a value between 1–8.

Syntax: [**no**] **tunnel source** *ipv4-address* | **ethernet** *port* | **loopback** *number* | **ve** *number*

The tunnel source can be an IP address or an interface.

For *ipv4-address*, use 8-bit values in dotted decimal notation.

The **ethernet** | **loopback** | **ve** parameter specifies an interface as the tunnel source. If you specify an Ethernet interface, also specify the port number associated with the interface. If you specify a loopback, VE, or interface, also specify the loopback, VE, or number, respectively.

Syntax: [**no**] **tunnel destination** *ipv4-address*

Specify the *ipv4-address* parameter using 8-bit values in dotted decimal notation.

Syntax: [**no**] **tunnel mode ipv6ip**

ipv6ip indicates that this is an IPv6 manual tunnel.

Syntax: ipv6 enable

The **ipv6 enable** command enables the tunnel. Alternatively, you could specify an IPv6 address, which would also enable the tunnel.

Syntax: ipv6 address *ipv6-prefix*/*prefix-length* [**eui-64**]

The **ipv6 address** command enables the tunnel. Alternatively, you could enter **ipv6 enable**, which would also enable the tunnel.

Specify the *ipv6-prefix* parameter in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

Specify the *prefix-length* parameter as a decimal value. A slash mark (/) must follow the *ipv6-prefix* parameter and precede the *prefix-length* parameter. The **eui-64** keyword configures the global address with an EUI-64 interface ID in the low-order 64 bits. The interface ID is automatically constructed in IEEE EUI-64 format using the interface's MAC address.

Clearing IPv6 tunnel statistics

You can clear statistics (reset all fields to zero) for all IPv6 tunnels or for a specific tunnel interface.

For example, to clear statistics for tunnel 1, enter the following command at the Privileged EXEC level or any of the Config levels of the CLI.

```
Brocade#clear ipv6 tunnel 1
```

To clear statistics for all IPv6 tunnels, enter the following command.

```
Brocade#clear ipv6 tunnel
```

Syntax: clear ipv6 tunnel [*number*]

The *number* parameter specifies the tunnel number.

Displaying IPv6 tunnel information

Use the commands in this section to display the configuration, status, and counters associated with IPv6 tunnels.

Displaying a summary of tunnel information

To display a summary of tunnel information, enter the following command at any level of the CLI.

```
Brocade#show ipv6 tunnel
IP6 Tunnels
  Tunnel  Mode         Packet Received  Packet Sent
  1       configured   0                0
  2       configured   0                22419
```

Syntax: show ipv6 tunnel

This display shows the following information.

TABLE 3 IPv6 tunnel summary information

Field	Description
Tunnel	The tunnel interface number.
Mode	The tunnel mode. Possible modes include the following: • configured – Indicates a manually configured tunnel.
Packet Received	The number of packets received by a tunnel interface. Note that this is the number of packets received by the CPU. It does not include the number of packets processed in hardware.
Packet Sent	The number of packets sent by a tunnel interface. Note that this is the number of packets sent by the CPU. It does not include the number of packets processed in hardware.

Displaying tunnel interface information

To display status and configuration information for tunnel interface 1, enter the following command at any level of the CLI.

```
Brocade#show interfaces tunnel 1
Tunnel1 is up, line protocol is up
  Hardware is Tunnel
  Tunnel source ve 30
  Tunnel destination is 10.2.2.10
  Tunnel mode ipv6ip
  No port name
  MTU 1480 bytes, encapsulation IPV4
```

Syntax: **show interfaces tunnel** *number*

The *number* parameter indicates the tunnel interface number for which you want to display information.

This display shows the following information.

TABLE 4 IPv6 tunnel interface information

Field	Description
Tunnel interface status	The status of the tunnel interface can be one of the following: • **up** – The tunnel mode is set and the tunnel interface is enabled. • **down** – The tunnel mode is not set. • **administratively down** – The tunnel interface was disabled with the **disable** command.
Line protocol status	The status of the line protocol can be one of the following: • **up** – IPv4 connectivity is established. • **down** – The line protocol is not functioning and is down.
Hardware is tunnel	The interface is a tunnel interface.
Tunnel source	The tunnel source can be one of the following: • An IPv4 address • The IPv4 address associated with an interface/port.
Tunnel destination	The tunnel destination can be an IPv4 address.

TABLE 4 IPv6 tunnel interface information (Continued)

Field	Description
Tunnel mode	The tunnel mode can be the following: • **ipv6ip** – indicates a manually configured tunnel
Port name	The port name configured for the tunnel interface.
MTU	The setting of the IPv6 maximum transmission unit (MTU).

Displaying interface level IPv6 settings

To display Interface level IPv6 settings for tunnel interface 1, enter the following command at any level of the CLI.

```
Brocade#show ipv6 inter tunnel 1
Interface Tunnel 1 is up, line protocol is up
  IPv6 is enabled, link-local address is fe80::3:4:2 [Preferred]
  Global unicast address(es):
    1001::1 [Preferred],  subnet is 1001::/64
    1011::1 [Preferred],  subnet is 1011::/64
  Joined group address(es):
    ff02::1:ff04:2
    ff02::5
    ff02::1:ff00:1
    ff02::2
    ff02::1
  MTU is 1480 bytes
  ICMP redirects are enabled
  No Inbound Access List Set
  No Outbound Access List Set
  OSPF enabled
```

The display command above reflects the following configuration.

```
Brocade#show running-config interface tunnel 1
!
interface tunnel 1
 port-name ManualTunnel1
 tunnel mode ipv6ip
 tunnel source loopback 1
 tunnel destination 10.1.1.1
 ipv6 address 1011::1/64
 ipv6 address 1001::1/64
 ipv6 ospf area 0
```

This display shows the following information.

TABLE 5 Interface level IPv6 tunnel information

Field	Description
Interface Tunnel status	The status of the tunnel interface can be one of the following: • **up** – IPv4 connectivity is established. • **down** – The tunnel mode is not set. • **administratively down** – The tunnel interface was disabled with the **disable** command.
Line protocol status	The status of the line protocol can be one of the following: • **up** – IPv6 is enabled through the **ipv6 enable** or **ipv6 address** command. • **down** – The line protocol is not functioning and is down.

ECMP load sharing for IPv6

The IPv6 route table selects the best route to a given destination from among the routes in the tables maintained by the configured routing protocols (BGP4, OSPF, static, and so on). The IPv6 route table can contain more than one path to a given destination. When this occurs, the Brocade device selects the path with the lowest cost for insertion into the routing table. If more than one path with the lowest cost exists, all of these paths are inserted into the routing table, subject to the configured maximum number of load sharing paths (by default 4). The device uses ***Equal-Cost Multi-Path (ECMP) load sharing*** to select a path to a destination.

When a route is installed by routing protocols or configured static route for the first time, and the IPv6 route table contains multiple, equal-cost paths to that route, the device checks the IPv6 neighbor for each next hop. Every next hop where the link layer address has been resolved will be stored in hardware. The device will initiate neighbor discovery for the next hops whose link layer addresses are not resolved. The hardware will hash the packet and choose one of the paths. The number of paths would be updated in hardware as the link layer gets resolved for a next hop.

If the path selected by the device becomes unavailable, the IPv6 neighbor should change state and trigger the update of the destination in the hardware.

Brocade FastIron devices support network-based ECMP load-sharing methods for IPv6 traffic. The Brocade device distributes traffic across equal-cost paths based on a XOR of some bits from the MAC source address, MAC destination address, IPv6 source address, IPv6 destination address, IPv6 flow label, IPv6 next header. The software selects a path based on a calculation involving the maximum number of load-sharing paths allowed and the actual number of paths to the destination network. This is the default ECMP load-sharing method for IPv6.

You can manually disable or enable ECMP load sharing for IPv6 and specify the number of equal-cost paths the device can distribute traffic across. In addition, you can display information about the status of ECMP load-sharing on the device.

Disabling or re-enabling ECMP load sharing for IPv6

ECMP load sharing for IPv6 is enabled by default. To disable the feature, enter the following command.

```
Brocade(config)#no ipv6 load-sharing
```

If you want to re-enable the feature after disabling it, you must specify the number of load-sharing paths. The maximum number of paths the device supports is a value from 2–8. By entering a command such as the following, iPv6 load-sharing will be re-enabled.

```
Brocade(config)#ipv6 load-sharing 4
```

Syntax: [**no**] **ipv6 load-sharing** *num*

The *num* parameter specifies the number of paths and can be from 2–8. The default is 4.

Changing the maximum load sharing paths for IPv6

By default, IPv6 ECMP load sharing allows traffic to be balanced across up to four equal paths. You can change the maximum number of paths the device supports to a value from 2–8.

To change the number of ECMP load sharing paths for IPv6, enter a command such as the following.

```
Brocade(config)#ipv6 load-sharing 6
```

Syntax: [**no**] **ipv6 load-sharing** [*num*]

The *num* parameter specifies the number of paths and can be from 2-8. The default is 4.

Enabling support for network-based ECMP load sharing for IPv6

Network-based ECMP load sharing is supported. In this configuration, traffic is distributed across equal-cost paths based on the destination network address. Routes to each network are stored in CAM and accessed when a path to a network is required. Because multiple hosts are likely to reside on a network, this method uses fewer CAM entries.

Displaying ECMP load-sharing information for IPv6

To display the status of ECMP load sharing for IPv6, enter the following command.

```
Brocade#show ipv6
Global Settings
  unicast-routing enabled, hop-limit 64
  No IPv6 Domain Name Set
  No IPv6 DNS Server Address set
  Prefix-based IPv6 Load-sharing is Enabled, Number of load share paths: 4
```

Syntax: **show ipv6**

Chapter

SNMP Access 6

Table 6 lists individual Brocade switches and the SNMP access methods they support. These features are supported in the Layer 2 and Layer 3 software images, except where explicitly noted.

TABLE 6 Supported SNMP access features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6430	ICX 6450
SNMP v1, v2, v3	Yes	Yes	Yes	Yes	Yes
Community strings	Yes	Yes	Yes	Yes	Yes
User-based security model for SNMP v3	Yes	Yes	Yes	Yes	Yes
SNMP v3 traps	Yes	Yes	Yes	Yes	Yes
Defining the UDP port for SNMP v3 traps	Yes	Yes	Yes	Yes	Yes
SNMP v3 over IPv6	Yes	Yes	Yes	Yes	Yes
AES encryption for SNMP v3	Yes	Yes	Yes	Yes	Yes

This chapter contains the following sections:

SNMP overview

SNMP is a set of protocols for managing complex networks. SNMP sends messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.

"Security Access" chapter in the *FastIron Ethernet Switch Security Configuration Guide* introduced a few methods used to secure SNMP access. They included the following:

- Using ACLs to restrict SNMP access
- Restricting SNMP access to a specific IP address
- Restricting SNMP access to a specific VLAN
- Disabling SNMP access

This chapter presents additional methods for securing SNMP access to Brocade devices. It contains the following sections:

- "SNMP community strings"
- "User-based security model"
- "SNMP v3 configuration examples"
- "SNMP version 3 traps"
- "Displaying SNMP Information"
- "SNMP v3 configuration examples"

Restricting SNMP access using ACL, VLAN, or a specific IP address constitute the first level of defense when the packet arrives at a Brocade device. The next level uses one of the following methods:

- Community string match In SNMP versions 1 and 2
- User-based model in SNMP version 3

SNMP views are incorporated in community strings and the user-based model.

SNMP community strings

SNMP versions 1 and 2 use community strings to restrict SNMP access.

- The default read-only community string is "public".
- There is no default read-write community string. You first must configure a read-write community string using the CLI. Then you can log on using "set" as the user name and the read-write community string you configure as the password.

You can configure as many additional read-only and read-write community strings as you need. The number of strings you can configure depends on the memory on the device. There is no practical limit.

NOTE
If you delete the startup-config file, the device automatically re-adds the default "public" read-only community string the next time you load the software.

Encryption of SNMP community strings

The software automatically encrypts SNMP community strings. Users with read-only access or who do not have access to management functions in the CLI cannot display the strings. For users with read-write access, the strings are encrypted in the CLI.

Encryption is enabled by default. You can disable encryption for individual strings or trap receivers if desired. Refer to the next section for information about encryption.

Adding an SNMP community string

The default SNMP community name (string) on a device is "public" with read only privilege.

You can assign other SNMP community strings, and indicate if the string is encrypted or clear. By default, the string is encrypted.

To add an encrypted community string, enter commands such as the following.

```
Brocade(config)#snmp-server community private rw
Brocade(config)#write memory
```

Syntax: **snmp-server community** [**0** | **1**] *string*
ro | **rw** [**view** *viewname*] [*standard-ACL-name* | *standard-ACL-id*]

The *string* parameter specifies the community string name. The string can be up to 32 characters long.

The **ro** | **rw** parameter specifies whether the string is r**ead-only (ro)** or **read-write (rw)**.

NOTE
If you issue a **no snmp-server community public ro** command and then enter a **write memory** command to save that configuration, the "public" community name is removed and will have no SNMP access. If for some reason the device is brought down and then brought up, the "no snmp-server community public ro" command is restored in the system and the "public" community string has no SNMP access.

The **0** | **1** parameter affects encryption for display of the string in the running-config and the startup-config file. Encryption is enabled by default. When encryption is enabled, the community string is encrypted in the CLI regardless of the access level you are using.

The encryption option can be omitted (the default) or can be one of the following:

- **0** – Disables encryption for the community string you specify with the command. The community string is shown as clear text in the running-config and the startup-config file. Use this option if you do not want the display of the community string to be encrypted.
- **1** – Assumes that the community string you enter is encrypted, and decrypts the value before using it.

NOTE
If you want the software to assume that the value you enter is the clear-text form, and to encrypt display of that form, do not enter **0** or **1**. Instead, omit the encryption option and allow the software to use the default behavior.

NOTE
If you specify encryption option **1**, the software assumes that you are entering the encrypted form of the community string. In this case, the software decrypts the community string you enter before using the value for authentication. If you accidentally enter option 1 followed by the clear-text version of the community string, authentication will fail because the value used by the software will not match the value you intended to use.

The command in the example above adds the read-write SNMP community string "private". When you save the new community string to the startup-config file (using the **write memory** command), the software adds the following command to the file.

```
snmp-server community 1 encrypted-string rw
```

To add a non-encrypted community string, you must explicitly specify that you do not want the software to encrypt the string. Here is an example.

```
Brocade(config)#snmp-server community 0 private rw
Brocade(config)#write memory
```

The command in this example adds the string "private" in the clear, which means the string is displayed in the clear. When you save the new community string to the startup-config file, the software adds the following command to the file.

```
snmp-server community 0 private rw
```

The **view** *viewname* parameter is optional. It allows you to associate a view to the members of this community string. Enter up to 32 alphanumeric characters. If no view is specified, access to the full MIB is granted. The view that you want must exist before you can associate it to a community string. Here is an example of how to use the view parameter in the community string command.

```
Brocade(config)#snmp-s community myread ro view sysview
```

The command in this example associates the view "sysview" to the community string named "myread". The community string has read-only access to "sysview". For information on how to create views, refer to "SNMP v3 configuration examples" on page 176.

The *standard-ACL-name* | *standard-ACL-id* parameter is optional. It allows you to specify which ACL group will be used to filter incoming SNMP packets. You can enter either the ACL name or its ID. Here are some examples.

```
Brocade(config)#snmp-s community myread ro view sysview 2
Brocade(config)#snmp-s community myread ro view sysview myACL
```

The command in the first example indicates that ACL group 2 will filter incoming SNMP packets; whereas, the command in the second example uses the ACL group called "myACL" to filter incoming packets.Refer to "Using ACLs to restrict SNMP access" section in the *FastIron Ethernet Switch Security Configuration Guide* for more information.

NOTE
To make configuration changes, including changes involving SNMP community strings, you must first configure a read-write community string using the CLI. Alternatively, you must configure another authentication method and log on to the CLI using a valid password for that method.

Displaying the SNMP community strings

To display the configured community strings, enter the following command at any CLI level.

```
Brocade#show snmp server
Contact: Marshall
Location: Copy Center
Community(ro): public
Community(rw): private
Traps
                   Cold start: Enable
                      Link up: Enable
                    Link down: Enable
               Authentication: Enable
     Locked address violation: Enable
         Power supply failure: Enable
                  Fan failure: Enable
          Temperature warning: Enable
                 STP new root: Enable
          STP topology change: Enable
                         ospf: Enable

 Total Trap-Receiver Entries: 4
Trap-Receiver IP Address       Community
     1        10.95.6.211
     2        10.95.5.21
```

Syntax: **show snmp server**

NOTE
If display of the strings is encrypted, the strings are not displayed. Encryption is enabled by default.

User-based security model

SNMP version 3 (RFC 2570 through 2575) introduces a User-Based Security model (RFC 2574) for authentication and privacy services.

SNMP version 1 and version 2 use community strings to authenticate SNMP access to management modules. This method can still be used for authentication. In SNMP version 3, the User-Based Security model of SNMP can be used to secure against the following threats:

- Modification of information
- Masquerading the identity of an authorized entity
- Message stream modification
- Disclosure of information

SNMP version 3 also supports View-Based Access Control Mechanism (RFC 2575) to control access at the PDU level. It defines mechanisms for determining whether or not access to a managed object in a local MIB by a remote principal should be allowed. (refer to “SNMP v3 configuration examples” on page 176.)

Configuring your NMS

In order to use the SNMP version 3 features.

1. Make sure that your Network Manager System (NMS) supports SNMP version 3.
2. Configure your NMS agent with the necessary users.
3. Configure the SNMP version 3 features in Brocade devices.

Configuring SNMP version 3 on Brocade devices

Follow the steps given below to configure SNMP version 3 on Brocade devices.

1. Enter an engine ID for the management module using the **snmp-server engineid** command if you will not use the default engine ID.Refer to "Defining the engine id" on page 166.
2. Create views that will be assigned to SNMP user groups using the **snmp-server view** command. refer to "SNMP v3 configuration examples" on page 176 for details.
3. Create ACL groups that will be assigned to SNMP user groups using the **access-list** command.
4. Create user groups using the **snmp-server group** command.Refer to "Defining an SNMP group" on page 167.
5. Create user accounts and associate these accounts to user groups using the **snmp-server user** command.Refer to "Defining an SNMP user account" on page 168.

If SNMP version 3 is not configured, then community strings by default are used to authenticate access.

Defining the engine id

A default engine ID is generated during system start up. To determine what the default engine ID of the device is, enter the **show snmp engineid** command and find the following line:

```
Local SNMP Engine ID: 800007c70300e05290ab60
```

See the section "Displaying the Engine ID" on page 174 for details.

The default engine ID guarantees the uniqueness of the engine ID for SNMP version 3. If you want to change the default engine ID, enter the **snmp-server engineid local** command.

```
Brocade(config)#snmp-server engineid local 800007c70300e05290ab60
```

Syntax: [**no**] **snmp-server engineid local** *hex-string*

The **local** parameter indicates that engine ID to be entered is the ID of this device, representing an SNMP management entity.

NOTE
Each user localized key depends on the SNMP server engine ID, so all users need to be reconfigured whenever the SNMP server engine ID changes.

NOTE
Since the current implementation of SNMP version 3 does not support Notification, remote engine IDs cannot be configured at this time.

The *hex-string* variable consists of 11 octets, entered as hexadecimal values. There are two hexadecimal characters in each octet. There should be an even number of hexadecimal characters in an engine ID.

The default engine ID has a maximum of 11 octets:

- Octets 1 through 4 represent the agent's SNMP management private enterprise number as assigned by the Internet Assigned Numbers Authority (IANA). The most significant bit of Octet 1 is "1". For example, “000007c7” is the ID for Brocade Communications, Inc. in hexadecimal. With Octet 1 always equal to "1", the first four octets in the default engine ID is always “800007c7” (which is 1991 in decimal).
- Octet 5 is always 03 in hexadecimal and indicates that the next set of values represent a MAC address.
- Octets 6 through 11 form the MAC address of the lowest port in the management module.

NOTE
Engine ID must be a unique number among the various SNMP engines in the management domain. Using the default engine ID ensures the uniqueness of the numbers.

Defining an SNMP group

SNMP groups map SNMP users to SNMP views. For each SNMP group, you can configure a read view, a write view, or both. Users who are mapped to a group will use its views for access control.

To configure an SNMP user group, enter a command such as the following.

```
Brocade(config)#snmp-server group admin v3 auth read all write all
```

Syntax: [**no**] **snmp-server group** *groupname* **v1** | **v2** | **v3** **auth** | **noauth** | **priv** [**access** *standard-ACL-id*] [**read** *viewstring* | **write** *viewstring*]

NOTE
This command is not used for SNMP version 1 and SNMP version 2. In these versions, groups and group views are created internally using community strings. (refer to “SNMP community strings” on page 162.) When a community string is created, two groups are created, based on the community string name. One group is for SNMP version 1 packets, while the other is for SNMP version 2 packets.

The **group** *groupname* parameter defines the name of the SNMP group to be created.

The **v1**, **v2**, or **v3** parameter indicates which version of SNMP is used. In most cases, you will be using v3, since groups are automatically created in SNMP versions 1 and 2 from community strings.

The **auth** | **noauth** parameter determines whether or not authentication will be required to access the supported views. If auth is selected, then only authenticated packets are allowed to access the view specified for the user group. Selecting **noauth** means that no authentication is required to access the specified view. Selecting **priv** means that an authentication password will be required from the users.

The **access** *standard-ACL-id* parameter is optional. It allows incoming SNMP packets to be filtered based on the standard ACL attached to the group.

The **read** *viewstring* | **write** *viewstring* parameter is optional. It indicates that users who belong to this group have either read or write access to the MIB.

The *viewstring* variable is the name of the view to which the SNMP group members have access. If no view is specified, then the group has no access to the MIB.

The value of *viewstring* is defined using the **snmp-server view** command. The SNMP agent comes with the "all" default view, which provides access to the entire MIB; however, it must be specified when creating the group. The "all" view also allows SNMP version 3 to be backwards compatibility with SNMP version 1 and version 2.

NOTE
If you will be using a view other than the "all" view, that view must be configured before creating the user group.Refer to the section "SNMP v3 configuration examples" on page 176, especially for details on the include | exclude parameters.

Defining an SNMP user account

The **snmp-server user** command does the following:

- Creates an SNMP user.
- Defines the group to which the user will be associated.
- Defines the type of authentication to be used for SNMP access by this user.
- Specifies one of the following encryption types used to encrypt the privacy password:
    - Data Encryption Standard (DES) – A symmetric-key algorithm that uses a 56-bit key.
    - Advanced Encryption Standard (AES) – The 128-bit encryption standard adopted by the U.S. government. This standard is a symmetric cipher algorithm chosen by the National Institute of Standards and Technology (NIST) as the replacement for DES.

Here is an example of how to create an SNMP User account.

```
Brocade(config)#snmp-s user bob admin v3 access 2 auth md5 bobmd5 priv des bobdes
```

The CLI for creating SNMP version 3 users has been updated as follows.

Syntax: [**no**] **snmp-server user** *name groupname* **v3**
[[**access** *standard-ACL-id*]
[[**encrypted**] [**auth md5** *md5-password* | **sha** *sha-password*]
[**priv** [**encrypted**] **des** *des-password-key* | **aes** *aes-password-key*]]]

The *name* parameter defines the SNMP user name or security name used to access the management module.

The *groupname* parameter identifies the SNMP group to which this user is associated or mapped. All users must be mapped to an SNMP group. Groups are defined using the **snmp-server group** command.

NOTE
The SNMP group to which the user account will be mapped should be configured before creating the user accounts; otherwise, the group will be created without any views. Also, ACL groups must be configured before configuring user accounts.

The **v3** parameter is required.

The **access** *standard-ACL-id* parameter is optional. It indicates that incoming SNMP packets are filtered based on the ACL attached to the user account.

NOTE
The ACL specified in a user account overrides the ACL assigned to the group to which the user is mapped. If no ACL is entered for the user account, then the ACL configured for the group will be used to filter packets.

The **encrypted** parameter means that the MD5 or SHA password will be a digest value. MD5 has 16 octets in the digest. SHA has 20. The digest string has to be entered as a hexadecimal string. In this case, the agent need not generate any explicit digest. If the **encrypted** parameter is not used, the user is expected to enter the authentication password string for MD5 or SHA. The agent will convert the password string to a digest, as described in RFC 2574.

The **auth md5 | sha** parameter is optional. It defines the type of encryption that the user must have to be authenticated. Choose between MD5 or SHA encryption. MD5 and SHA are two authentication protocols used in SNMP version 3.

The *md5-password* and *sha-password* define the password the user must use to be authenticated. These password must have a minimum of 8 characters. If the encrypted parameter is used, then the digest has 16 octets for MD5 or 20 octets for SHA.

NOTE
Once a password string is entered, the generated configuration displays the digest (for security reasons), not the actual password.

The **priv** [**encrypted**] parameter is optional after you enter the md5 or sha password. The **priv** parameter specifies the encryption type (DES or AES) used to encrypt the privacy password. If the **encrypted** keyword is used, do the following:

- If DES is the privacy protocol to be used, enter **des** followed by a 16-octet DES key in hexadecimal format for the *des-password-key*. If you include the encrypted keyword, enter a password string of at least 8 characters.
- If AES is the privacy protocol to be used, enter **aes** followed by the AES password key. For a small password key, enter 12 characters. For a big password key, enter 16 characters. If you include the encrypted keyword, enter a password string containing 32 hexadecimal characters.

Defining SNMP views

SNMP views are named groups of MIB objects that can be associated with user accounts to allow limited access for viewing and modification of SNMP statistics and system configuration. SNMP views can also be used with other commands that take SNMP views as an argument. SNMP views reference MIB objects using object names, numbers, wildcards, or a combination of the three. The numbers represent the hierarchical location of the object in the MIB tree. You can reference individual objects in the MIB tree or a subset of objects from the MIB tree.

To configure the number of SNMP views available on the Brocade device, enter the following command.

```
Brocade(config)#system-max view 15
```

Syntax: system-max view *number-of-views*

This command specifies the maximum number of SNMPv2 and v3 views that can be configured on a device. The number of views can be from 10 – 65536. The default is 10 views.

To add an SNMP view, enter one of the following commands.

```
Brocade(config)#snmp-server view Maynes system included
Brocade(config)#snmp-server view Maynes system.2 excluded
Brocade(config)#snmp-server view Maynes 2.3.*.6 included
Brocade(config)#write mem
```

NOTE
The **snmp-server view** command supports the MIB objects as defined in RFC 1445.

Syntax: [**no**] **snmp-server view** *name mib_tree* **included** | **excluded**

The *name* parameter can be any alphanumeric name you choose to identify the view. The names cannot contain spaces.

The *mib_tree* parameter is the name of the MIB object or family. MIB objects and MIB sub-trees can be identified by a name or by the numbers called Object Identifiers (OIDs) that represent the position of the object or sub-tree in the MIB hierarchy. You can use a wildcard (*) in the numbers to specify a sub-tree family.

The **included** | **excluded** parameter specifies whether the MIB objects identified by the *mib_family* parameter are included in the view or excluded from the view.

NOTE
All MIB objects are automatically excluded from any view unless they are explicitly included; therefore, when creating views using the **snmp-server view** command, indicate which portion of the MIB you want users to access.

For example, you may want to assign the view called “admin” a community string or user group. The “admin” view will allow access to the Brocade MIBs objects that begin with the 1.3.6.1.4.1.1991 object identifier. Enter the following command.

```
Brocade(config)#snmp-server view admin 1.3.6.1.4.1.1991 included
```

You can exclude portions of the MIB within an inclusion scope. For example, if you want to exclude the snAgentSys objects, which begin with 1.3.6.1.4.1.1991.1.1.2 object identifier from the admin view, enter a second command such as the following.

```
Brocade(config)#snmp-server view admin 1.3.6.1.4.1.1991.1.1.2 excluded
```

NOTE
Note that the exclusion is within the scope of the inclusion.

To delete a view, use the no parameter before the command.

SNMP version 3 traps

Brocade devices support SNMP notifications in SMIv2 format. This allows notifications to be encrypted and sent to the target hosts in a secure manner.

Defining an SNMP group and specifying which view is notified of traps

The SNMP group command allows configuration of a viewname for notification purpose, similar to the read and write view. The default viewname is "all", which allows access to the entire MIB.

To configure an SNMP user group, first configure SNMPv3 views using the **snmp-server view** command. Refer to "SNMP v3 configuration examples" on page 176. Then enter a command such as the following.

```
Brocade(config)#snmp-server group admin v3 auth read all write all
notify all
```

Syntax: [**no**] **snmp-server group** *groupname*
v1 | **v2** | **v3**
auth | **noauth** | **priv**
[**access** *standard-ACL-id*] [**read** *viewstring* | **write** *viewstring* | **notify** *viewstring*]

The **group** *groupname* parameter defines the name of the SNMP group to be created.

The **v1**, **v2**, or **v3** parameter indicates which version of SNMP to use. In most cases, you will use v3, since groups are automatically created in SNMP versions 1 and 2 from community strings.

The **auth** | **noauth** parameter determines whether or not authentication will be required to access the supported views. If auth is selected, then only authenticated packets are allowed to access the view specified for the user group. Selecting **noauth** means that no authentication is required to access the specified view. Selecting **priv** means that an authentication password will be required from the users.

The **access** *standard-ACL-id* parameter is optional. It allows incoming SNMP packets to be filtered based on the standard ACL attached to the group.

The **read** *viewstring* | **write** *viewstring* parameter is optional. It indicates that users who belong to this group have either read or write access to the MIB.

The **notify** view allows administrators to restrict the scope of varbind objects that will be part of the notification. All of the varbinds need to be in the included view for the notification to be created.

The *viewstring* variable is the name of the view to which the SNMP group members have access. If no view is specified, then the group has no access to the MIB.

Defining the UDP port for SNMP v3 traps

The SNMP host command enhancements allow configuration of notifications in SMIv2 format, with or without encryption, in addition to the previously supported SMIv1 trap format.

You can define a port that receives the SNMP v3 traps by entering a command such as the following.

```
Brocade(config)#snmp-server host 192.168.4.11 version v3 auth security-name port
4/1
```

Syntax: [**no**] **snmp-server host** *ip-addr* | *ipv6-addr* **version** [**v1** | **v2c** *community-string* | **v3 auth** | **noauth** | **priv** *security-name*] [**port** *trap-UDP-port-number*]

The *ip-addr* parameter specifies the IP address of the host that will receive the trap.

For **version**, indicate one of the following

For SNMP version 1, enter **v1** and the name of the community string (*community-string*). This string is encrypted within the system.

NOTE
If the configured version is v2c, then the notification is sent out in SMIv2 format, using the community string, but in cleartext mode. To send the SMIv2 notification in SNMPv3 packet format, configure v3 with auth or privacy parameters, or both, by specifying a security name. The actual authorization and privacy values are obtained from the security name.

For SNMP version 2c, enter **v2** and the name of the community string. This string is encrypted within the system.

For SNMP version 3, enter one of the following depending on the authorization required for the host:

- **v3 auth** *security-name*: Allow only authenticated packets.
- **v3 no auth** *security-name*: Allow all packets.
- **v3 priv** *security-name*: A password is required

For **port** *trap-UDP-port-number*, specify the UDP port number on the host that will receive the trap.

Trap MIB changes

To support the SNMP V3 trap feature, the Brocade Enterprise Trap MIB was rewritten in SMIv2 format, as follows:

- The MIB name was changed from FOUNDRY-SN-TRAP-MIB to FOUNDRY-SN-NOTIFICATION-MIB
- Individual notifications were changed to NOTIFICATION-TYPE instead of TRAP-TYPE.
- As per the SMIv2 format, each notification has an OID associated with it. The root node of the notification is snTraps (OID enterprise.foundry.0). For example, OID for snTrapRunningConfigChanged is {snTraps.73}. Earlier, each trap had a trap ID associated with it, as per the SMIv1 format.

Backward compatibility with SMIv1 trap format

The Brocade device will continue to support creation of traps in SMIv1 format, as before. To allow the device to send notifications in SMIv2 format, configure the device as described above. The default mode is still the original SMIv1 format.

Specifying an IPv6 host as an SNMP trap receiver

You can specify an IPv6 host as a trap receiver to ensure that all SNMP traps sent by the device will go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. To do so, enter a command such as the following.

```
Brocade(config)#snmp-server host ipv6 2001:DB8:89::13
```

Syntax: **snmp-server host ipv6** *ipv6-address*

The *ipv6-address* must be in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

SNMP v3 over IPv6

Some FastIron devices support IPv6 for SNMP version 3.

Restricting SNMP Access to an IPv6 Node

You can restrict SNMP access so that the Brocade device can only be accessed by the IPv6 host address that you specify. To do so, enter a command such as the following .

```
Brocade(config)#snmp-client ipv6 2001:DB8:89::23
```

Syntax: **snmp-client ipv6** *ipv6-address*

The *ipv6-address* must be in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

Specifying an IPv6 host as an SNMP trap receiver

You can specify an IPv6 host as a trap receiver to ensure that all SNMP traps sent by the Brocade device will go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. To do so, enter the **snmp-server host ipv6** command .

```
Brocade(config)#snmp-server host ipv6 2001:DB8:89::13
```

Syntax: **snmp-server host ipv6** *ipv6-address*

The *ipv6-address* must be in hexadecimal format using 16-bit values between colons as documented in RFC 2373.

Viewing IPv6 SNMP server addresses

Many of the existing **show** commands display IPv6 addresses for IPv6 SNMP servers. The following example shows output for the **show snmp server** command.

```
Brocade#show snmp server
      Contact:
     Location:
Community(ro): .....
Traps
              Warm/Cold start: Enable
                      Link up: Enable
                    Link down: Enable
               Authentication: Enable
     Locked address violation: Enable
         Power supply failure: Enable
                  Fan failure: Enable
          Temperature warning: Enable
                 STP new root: Enable
          STP topology change: Enable
                         vsrp: Enable

  Total Trap-Receiver Entries: 4
Trap-Receiver IP-Address                Port-Number Community
     1        10.147.201.100               162     .....
     2        2001:DB8::200                162     .....
      3        10.147.202.100               162     .....
      4       2001:DB8::200                162     .....
```

Displaying SNMP Information

This section lists the commands for viewing SNMP-related information.

Displaying the Engine ID

To display the engine ID of a management module, enter a command such as the following.

```
Brocade#show snmp engineid
Local SNMP Engine ID: 800007c70300e05290ab60
Engine Boots: 3
Engine time: 5
```

Syntax: **show snmp engineid**

The engine ID identifies the source or destination of the packet.

The engine boots represents the number of times that the SNMP engine reinitialized itself with the same engine ID. If the engineID is modified, the boot count is reset to 0.

The engine time represents the current time with the SNMP agent.

Displaying SNMP groups

To display the definition of an SNMP group, enter a command such as the following.

```
Brocade#show snmp group
groupname = exceptifgrp
security model = v3
security level = authNoPriv
ACL id = 2
readview = exceptif
writeview = none
```

Syntax: **show snmp group**

The value for security level can be one of the following.

Security level	Authentication
none	If the security model shows v1 or v2, then security level is blank. User names are not used to authenticate users; community strings are used instead.
noauthNoPriv	Displays if the security model shows v3 and user authentication is by user name only.
authNoPriv	Displays if the security model shows v3 and user authentication is by user name and the MD5 or SHA algorithm.

Displaying user information

To display the definition of an SNMP user account, enter a command such as the following.

```
Brocade#show snmp user
username = bob
ACL id = 2
group = admin
security model = v3
group ACL id = 0
authtype = md5
authkey = 3aca18d90b8d172760e2dd2e8f59b7fe
privtype = des,  privkey = 1088359afb3701730173a6332d406eec
engine ID= 800007c70300e052ab0000
```

Syntax: **show snmp user**

Interpreting varbinds in report packets

If an SNMP version 3 request packet is to be rejected by an SNMP agent, the agent sends a report packet that contains one or more varbinds. The varbinds contain additional information, showing the cause of failures. An SNMP manager application decodes the description from the varbind. The following table presents a list of varbinds supported by the SNMP agent.

Varbind object Identifier	Description
1. 3. 6. 1. 6. 3. 11. 2. 1. 3. 0	Unknown packet data unit.
1. 3. 6. 1. 6. 3. 12. 1. 5. 0	The value of the varbind shows the engine ID that needs to be used in the snmp-server engineid command
1. 3. 6. 1. 6. 3. 15. 1. 1. 1. 0	Unsupported security level.

Varbind object Identifier	Description
1. 3. 6. 1. 6. 3. 15. 1. 1. 2. 0	Not in time packet.
1. 3. 6. 1. 6. 3. 15. 1. 1. 3. 0	Unknown user name. This varbind may also be generated: • If the configured ACL for this user filters out this packet. • If the group associated with the user is unknown.
1. 3. 6. 1. 6. 3. 15. 1. 1. 4. 0	Unknown engine ID. The value of this varbind would be the correct authoritative engineID that should be used.
1. 3. 6. 1. 6. 3. 15. 1. 1. 5. 0	Wrong digest.
1. 3. 6. 1. 6. 3. 15. 1. 1. 6. 0	Decryption error.

SNMP v3 configuration examples

The following sections present examples of how to configure SNMP v3.

Simple SNMP v3 configuration

```
Brocade(config)#snmp-s group admingrp v3 priv read all write all notify all
Brocade(config)#snmp-s user adminuser admingrp v3 auth md5 auth password priv
privacy password
Brocade(config)#snmp-s host dest-ip version v3 privacy adminuser
```

More detailed SNMP v3 configuration

```
Brocade(config)#snmp-server view internet internet included
Brocade(config)#snmp-server view system system included
Brocade(config)#snmp-server community ..... ro
Brocade(config)#snmp-server community ..... rw
Brocade(config)#snmp-server contact isc-operations
Brocade(config)#snmp-server location sdh-pillbox
Brocade(config)#snmp-server host 128.91.255.32 .....
Brocade(config)#snmp-server group ops v3 priv read internet write system
Brocade(config)#snmp-server group admin v3 priv read internet write internet
Brocade(config)#snmp-server group restricted v3 priv read internet
Brocade(config)#snmp-server user ops ops v3 encrypted auth md5
ab8e9cd6d46e7a270b8c9549d92a069 priv encrypted des
0e1b153303b6188089411447dbc32de
Brocade(config)#snmp-server user admin admin v3 encrypted auth md5
0d8a2123f91bfbd8695fef16a6f4207b priv encrypted des
18e0cf359fce4fcd60df19c2b6515448
Brocade(config)#snmp-server user restricted restricted v3 encrypted auth md5
261fd8f56a3ad51c8bcec1e4609f54dc priv encrypted des
d32e66152f89de9b2e0cb17a65595f43
```

Chapter

7

Foundry Discovery Protocol (FDP) and Cisco Discovery Protocol (CDP) Packets

Table 7 lists individual Brocade switches and the discovery protocols they support. These features are supported in the Layer 2 and Layer 3 software images, except where explicitly noted.

TABLE 7 Supported discovery protocol features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6430	ICX 6450
Foundry Discovery Protocol (FDP) for IPv4 and IPv6 traffic	Yes	Yes	Yes	Yes	Yes
Cisco Discovery Protocol (CDP) for IPv4 and IPV6 traffic	Yes	Yes	Yes	Yes	Yes

This chapter contains the following sections:

FDP Overview

The Foundry Discovery Protocol (FDP) enables Brocade devices to advertise themselves to other Brocade devices on the network. When you enable FDP on a Brocade device, the device periodically advertises information including the following:

- Hostname (device ID)
- Product platform and capability
- Software version
- VLAN and Layer 3 protocol address information for the port sending the update. IP, IPX, and AppleTalk Layer 3 information is supported.

A Brocade device running FDP sends FDP updates on Layer 2 to MAC address 00-00-00-CC-CC-CC. Other Brocade devices listening on that address receive the updates and can display the information in the updates. Brocade devices can send and receive FDP updates on Ethernet interfaces.

FDP is disabled by default.

NOTE
If FDP is not enabled on a Brocade device that receives an FDP update or the device is running a software release that does not support FDP, the update passes through the device at Layer 2.

FDP configuration

The following sections describe how to enable Foundry Discovery Protocol (FDP) and how to change the FDP update and hold timers.

Enabling FDP globally

To enable a Brocade device to globally send FDP packets, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)# fdp run
```

Syntax: [no] fdp run

The feature is disabled by default.

Enabling FDP at the interface level

By default, FDP is enabled at the interface level after FDP is enabled on the device.

When FDP is enabled globally, you can disable and re-enable FDP on individual ports.

Disable FDP by entering commands such as the following:

```
Brocade(config)# int e 2/1
Brocade(config-if-2/1)# no fdp enable
```

Enable or repenable FDP by entering commands such as the following:

```
Brocade(config-if-2/1)# fdp enable
```

Syntax: [no] fdp enable

Specifying the IP management address to advertise

When FDP is enabled, by default, the Brocade device advertises one IPv4 address and one IPv6 address to its FDP neighbors. If desired, you can configure the device to advertise only the IPv4 management address or only the IPv6 management address. You can set the configuration globally on a Layer 2 switch, or on an interface on a Layer 3 switch.

For example, to configure a Layer 2 switch to advertise the IPv4 address, enter the following command at the Global CONFIG level of the CLI:

```
Brocade(config)# fdp advertise ipv4
```

To configure a Layer 3 switch to advertise the IPv6 address, enter the following command at the Interface level of the CLI:

```
Brocade(config-if-2/1)# fdp advertise ipv6
```

Syntax: fdp advertise ipv4 | ipv6

Changing the FDP update timer

By default, a Brocade device enabled for FDP sends an FDP update every 60 seconds. You can change the update timer to a value from 5 – 900 seconds.

To change the FDP update timer, enter a command such as the following at the global CONFIG level of the CLI.

```
Brocade(config)# fdp timer 120
```

Syntax: [**no**] **fdp timer** *secs*

The *secs* parameter specifies the number of seconds between updates and can be from 5 – 900 seconds. The default is 60 seconds.

Changing the FDP hold time

By default, a Brocade device that receives an FDP update holds the information until one of the following events occurs:

- The device receives a new update.
- 180 seconds have passed since receipt of the last update. This is the hold time.

Once either of these events occurs, the device discards the update.

To change the FDP hold time, enter the **fdp holdtime** command at the global CONFIG level of the CLI.

```
Brocade(config)# fdp holdtime 360
```

Syntax: [**no**] **fdp holdtime** *secs*

The *secs* parameter specifies the number of seconds a Brocade device that receives an FDP update can hold the update before discarding it. You can specify from 10 – 255 seconds. The default is 180 seconds.

Displaying FDP information

You can display the following Foundry Discovery Protocol (FDP) information:

- FDP entries for Brocade neighbors
- Individual FDP entries
- FDP information for an interface on the device you are managing
- FDP packet statistics

NOTE
If the Brocade device has intercepted CDP updates, then the CDP information is also displayed.

Displaying neighbor information

To display a summary list of all the Brocade neighbors that have sent FDP updates to this Brocade device, enter the **show fdp neighbor** command.

```
Brocade# show fdp neighbor
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
(*) indicates a CDP device

   Device ID      Local Int    Holdtm Capability Platform     Port ID
   -------------- ------------ ------ ---------- -----------  -------------
   FastIronB       Eth 2/9     178    Router     FastIron Rou Eth 2/9
```

Syntax: **show fdp neighbor** [**ethernet** *port*] [**detail**]

The **ethernet** *port* parameter lists the information for updates received on the specified port.

The **detail** parameter lists detailed information for each device.

The **show fdp neighbor** command, without optional parameters, displays the following information.

TABLE 8 Summary FDP and CDP neighbor information

This line...	Displays...
Device ID	The hostname of the neighbor.
Local Int	The interface on which this Brocade device received an FDP or CDP update for the neighbor.
Holdtm	The maximum number of seconds this device can keep the information received in the update before discarding it.
Capability	The role the neighbor is capable of playing in the network.
Platform	The product platform of the neighbor.
Port ID	The interface through which the neighbor sent the update.

To display detailed information, enter the **show fdp neighbor detail** command.

```
BrocadeA# show fdp neighbor detail
Device ID: FastIronB configured as default VLAN1, tag-type8100
Entry address(es):
  IP address: 192.168.0.13
  IPv6 address (Global): c:a:f:e:c:a:f:e
Platform: FastIron Router,  Capabilities: Router
Interface: Eth 2/9
Port ID (outgoing port): Eth 2/9 is TAGGED in following VLAN(s):
 9 10 11
Holdtime : 176 seconds
Version :
Foundry, Inc. Router, IronWare Version 07.6.01b1T53 Compiled on Aug 29
2002 at 10:35:21 labeled as B2R07601b1
```

The **show fdp neighbor detail** command displays the following information.

TABLE 9 Detailed FDP and CDP neighbor information

Parameter	Definition
Device ID	The hostname of the neighbor. In addition, this line lists the VLAN memberships and other VLAN information for the neighbor port that sent the update to this device.
Entry address(es)	The Layer 3 protocol addresses configured on the neighbor port that sent the update to this device. If the neighbor is a Layer 2 Switch, this field lists the management IP address.
Platform	The product platform of the neighbor.
Capabilities	The role the neighbor is capable of playing in the network.
Interface	The interface on which this device received an FDP or CDP update for the neighbor.
Port ID	The interface through which the neighbor sent the update.
Holdtime	The maximum number of seconds this device can keep the information received in the update before discarding it.
Version	The software version running on the neighbor.

Displaying FDP entries

To display the detailed neighbor information for a specific device, enter the **show fdp entry FastIron***x* command.

```
BrocadeA# show fdp entry FastIronB
Device ID: FastIronB configured as default VLAN1, tag-type8100
Entry address(es):
Platform: FastIron Router,  Capabilities: Router
Interface: Eth 2/9
Port ID (outgoing port): Eth 2/9 is TAGGED in following VLAN(s):
 9 10 11
Holdtime : 176 seconds
Version :
Foundry, Inc. Router, IronWare Version 07.6.01b1T53 Compiled on Aug 29
2002 at 10:35:21 labeled as B2R07601b1
```

Syntax: **show fdp entry** * | *device-id*

The * | *device-id* parameter specifies the device ID. If you enter *, the detailed updates for all neighbor devices are displayed. If you enter a specific device ID, the update for that device is displayed. For information about the display, refer to Table 9.

Displaying FDP information for an interface

To display FDP information for an interface, enter a command such as the following.

```
BrocadeA# show fdp interface ethernet 2/3
FastEthernet2/3 is up, line protocol is up
  Encapsulation ethernet
  Sending FDP packets every 5 seconds
  Holdtime is 180 seconds
```

This example shows information for Ethernet port 2/3. The port sends FDP updates every 5 seconds. Neighbors that receive the updates can hold them for up to 180 seconds before discarding them.

Syntax: **show fdp interface** [**ethernet** *port*]

The **ethernet** *port* parameter lists the information only for the specified interface.

Displaying FDP and CDP statistics

To display FDP and CDP packet statistics, enter the following command.

```
BrocadeA# show fdp traffic
CDP/FDP counters:
  Total packets output: 6, Input: 5
  Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
  No memory: 0, Invalid packet: 0, Fragmented: 0
  Internal errors: 0
```

Syntax: **show fdp traffic**

Clearing FDP and CDP information

You can clear the following FDP and CDP information:

- Information received in FDP and CDP updates
- FDP and CDP statistics

The same commands clear information for both FDP and CDP.

Clearing FDP and CDP neighbor information

To clear the information received in FDP and CDP updates from neighboring devices, enter the following command.

```
Brocade# clear fdp table
```

Syntax: **clear fdp table**

NOTE
This command clears all the updates for FDP and CDP.

Clearing FDP and CDP statistics

To clear FDP and CDP statistics, enter the following command.

```
Brocade# clear fdp counters
```

Syntax: **clear fdp counters**

CDP packets

Cisco Discovery Protocol (CDP) packets are used by Cisco devices to advertise themselves to other Cisco devices. By default, Brocade devices forward these packets without examining their contents. You can configure a Brocade device to intercept and display the contents of CDP packets. This feature is useful for learning device and interface information for Cisco devices in the network.

Brocade devices support intercepting and interpreting CDP version 1 and version 2 packets.

NOTE
The Brocade device can interpret only the information fields that are common to both CDP version 1 and CDP version 2.

NOTE
When you enable interception of CDP packets, the Brocade device drops the packets. As a result, Cisco devices will no longer receive the packets.

Enabling interception of CDP packets globally

To enable the device to intercept and display CDP packets, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)# cdp run
```

Syntax: [**no**] **cdp run**

The feature is disabled by default.

Enabling interception of CDP packets on an interface

You can disable and enable CDP at the interface level.

You can enter commands such as the following.

```
Brocade(config)# int e 2/1
Brocade(config-if-2/1)# cdp enable
```

Syntax: [**no**] **cdp enable**

By default, the feature is enabled on an interface once CDP is enabled on the device.

Displaying CDP information

You can display the following CDP information:

- Cisco neighbors
- CDP entries for all Cisco neighbors or a specific neighbor
- CDP packet statistics

Displaying neighbors

To display the Cisco neighbors the Brocade device has learned from CDP packets, enter the **show fdp neighbors** command.

```
Brocade# show fdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
(*) indicates a Cisco device

   Device ID      Local Int    Holdtm Capability Platform    Port ID
   -------------- ------------ ------ ---------- ----------- -------------
(*)Router         Eth 1/1      124    R          cisco RSP4
FastEthernet5/0/0
```

To display detailed information for the neighbors, enter the **show fdp neighbors detail** command.

```
Brocade# show fdp neighbors detail
Device ID: Router
Entry address(es):
  IP address: 10.95.6.143
Platform: cisco RSP4,  Capabilities: Router
Interface: Eth 1/1,  Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 150 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1,  RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

To display information about a neighbor attached to a specific port, enter a command such as the following.

```
Brocade# show fdp neighbors ethernet 1/1
Device ID: Router
Entry address(es):
  IP address: 10.95.6.143
Platform: cisco RSP4,  Capabilities: Router
Interface: Eth 1/1,  Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 127 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1,  RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

Syntax: **show fdp neighbors** [**detail** | **ethernet** *port*]

Displaying CDP entries

To display CDP entries for all neighbors, enter the **show fdp entry** command.

```
Brocade# show fdp entry *
Device ID: Router
Entry address(es):
  IP address: 10.95.6.143
Platform: cisco RSP4,  Capabilities: Router
Interface: Eth 1/1,  Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 124 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1,  RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

To display CDP entries for a specific device, specify the device ID, as shown in the following example.

```
Brocade# show fdp entry Router1
Device ID: Router1
Entry address(es):
  IP address: 10.95.6.143
Platform: cisco RSP4,  Capabilities: Router
Interface: Eth 1/1,  Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 156 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1,  RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

Syntax: **show fdp entry** * | *device-id*

Displaying CDP statistics

To display CDP packet statistics, enter the **show fdp traffic** command.

```
Brocade# show fdp traffic
CDP counters:
  Total packets output: 0, Input: 3
  Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
  No memory: 0, Invalid packet: 0, Fragmented: 0
```

Syntax: **show fdp traffic**

Clearing CDP information

You can clear the following CDP information:

- Cisco Neighbor information
- CDP statistics

To clear the Cisco neighbor information, enter the **clear fdp table** command.

```
Brocade# clear fdp table
```

Syntax: clear fdp table

To clear CDP statistics, enter the following command.

```
Brocade# clear fdp counters
```

Syntax: clear fdp counters

Chapter

LLDP and LLDP-MED

8

Table 10 lists the individual Brocade FastIron switches and the Link Layer Discovery Protocol (LLDP) features they support. These features are supported in the Layer 2 and Layer 3 software images, except where explicitly noted.

TABLE 10 Supported LLDP features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6430	ICX 6450
LLDP	Yes	Yes	Yes	Yes	Yes
LLDP-MED	Yes	Yes	Yes	Yes	Yes
Support for tagged LLDP packets	Yes	Yes	Yes	Yes	Yes
IPv4 management address advertisement	Yes	Yes	Yes	Yes	Yes
IPv6 management address advertisement	Yes	Yes	Yes	Yes	Yes
LLDP operating mode setting per port	Yes	Yes	Yes	Yes	Yes
LLDP processing on 802.1x blocked port	Yes	Yes	Yes	Yes	Yes
Setting the maximum number of LLDP neighbors	Yes	Yes	Yes	Yes	Yes
SNMP and Syslog messages	Yes	Yes	Yes	Yes	Yes
LLDP transmission intervals	Yes	Yes	Yes	Yes	Yes
Holdtime multiplier for transmit TTL	Yes	Yes	Yes	Yes	Yes
Configuring the minimum time between port reinitializations	Yes	Yes	Yes	Yes	Yes
Fast start repeat count for LLDP-MED	Yes	Yes	Yes	Yes	Yes
Location ID for LLDP-MED	Yes	Yes	Yes	Yes	Yes
LLDP-MED network policy	Yes	Yes	Yes	Yes	Yes
LLDP statistics and configuration details	Yes	Yes	Yes	Yes	Yes

This chapter contains the following sections:

This chapter describes how to configure the following protocols:

Link layer discovery protocol (LLDP) – The Layer 2 network discovery protocol described in the IEEE 802.1AB standard, *Station and Media Access Control Connectivity Discovery*. This protocol enables a station to advertise its capabilities to, and to discover, other LLDP-enabled stations in the same 802 LAN segments.

LLDP media endpoint devices (LLDP-MED) – The Layer 2 network discovery protocol extension described in the ANSI/TIA-1057 standard, *LLDP for Media Endpoint Devices*. This protocol enables a switch to configure and manage connected Media Endpoint devices that need to send media streams across the network (e.g., IP telephones and security cameras).

LLDP enables network discovery between Network Connectivity devices (such as switches), whereas LLDP-MED enables network discovery at the edge of the network, between Network Connectivity devices and media Endpoint devices (such as IP phones).

The information generated by LLDP and LLDP-MED can be used to diagnose and troubleshoot misconfigurations on both sides of a link. For example, the information generated can be used to discover devices with misconfigured or unreachable IP addresses, and to detect port speed and duplex mismatches.

LLDP and LLDP-MED facilitate interoperability across multiple vendor devices. Brocade devices running LLDP can interoperate with third-party devices running LLDP.

The Brocade LLDP and LLDP-MED implementation adheres to the IEEE 802.1AB and TIA-1057 standards.

LLDP terms used in this chapter

Endpoint device – An LLDP-MED device located at the network edge, that provides some aspect of IP communications service based on IEEE 802 LAN technology. An Endpoint device is classified in one of three class types (I, II, or III) and can be an IP telephone, softphone, VoIP gateway, or conference bridge, among others.

LLDP agent – The protocol entity that implements LLDP for a particular IEEE 802 device. Depending on the configured LLDP operating mode, an LLDP agent can send and receive LLDP advertisements (frames), or send LLDP advertisements only, or receive LLDP advertisements only.

LLDPDU (LLDP Data Unit) – A unit of information in an LLDP packet that consists of a sequence of short variable length information elements, known as **TLVs**. LLDP pass-through is not supported in conformance to IEEE standard.

MIB (Management Information Base) – A virtual database that identifies each manageable object by its name, syntax, accessibility, and status, along with a text description and unique object identifier (OID). The database is accessible by a Network Management Station (NMS) using a management protocol such as the Simple Network Management Protocol (SNMP).

Network connectivity device – A forwarding 802 LAN device, such as a router, switch, or wireless access point.

Station – A node in a network.

TLV (Type-Length-Value) – An information element in an LLDPDU that describes the type of information being sent, the length of the information string, and the value (actual information) that will be transmitted.

TTL (Time-to-Live) – Specifies the length of time that the receiving device should maintain the information acquired through LLDP in its MIB.

LLDP overview

LLDP enables a station attached to an IEEE 802 LAN/MAN to advertise its capabilities to, and to discover, other stations in the same 802 LAN segments.

The information distributed by LLDP (the advertisement) is stored by the receiving device in a standard Management Information Base (MIB), accessible by a Network Management System (NMS) using a management protocol such as the Simple Network Management Protocol (SNMP). The information also can be viewed from the CLI, using **show LLDP** commands.

Figure 2 illustrates LLDP connectivity

FIGURE 2 **LLDP connectivity**

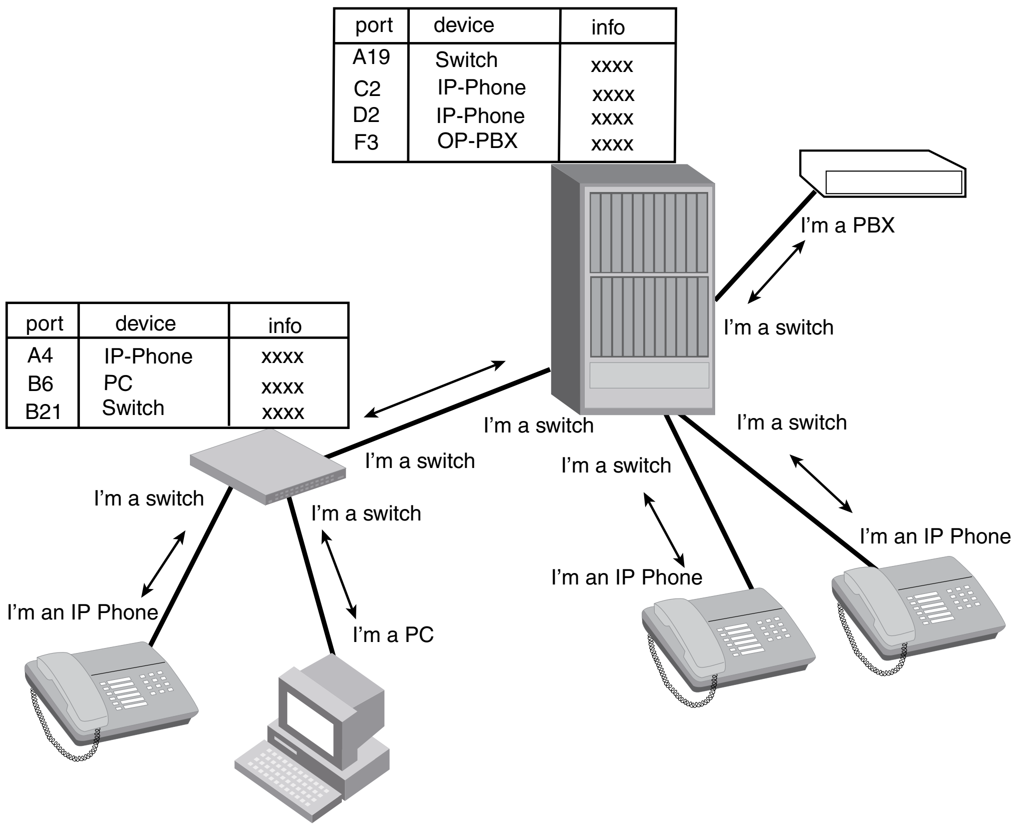

Benefits of LLDP

LLDP provides the following benefits:

- Network Management:
  - Simplifies the use of and enhances the ability of network management tools in multi-vendor environments
  - Enables discovery of accurate physical network topologies such as which devices are neighbors and through which ports they connect
  - Enables discovery of stations in multi-vendor environments
- Network Inventory Data:
  - Supports optional system name, system description, system capabilities and management address
  - System description can contain the device product name or model number, version of hardware type, and operating system
  - Provides device capability, such as switch, router, or WLAN access point
- Network troubleshooting:
  - Information generated by LLDP can be used to detect speed and duplex mismatches
  - Accurate topologies simplify troubleshooting within enterprise networks
  - Can discover devices with misconfigured or unreachable IP addresses

LLDP-MED overview

LLDP-MED is an extension to LLDP. This protocol enables advanced LLDP features in a Voice over IP (VoIP) network. Whereas LLDP enables network discovery between Network Connectivity devices, LLDP-MED enables network discovery between Network Connectivity devices and media Endpoints such as, IP telephones, softphones, VoIP gateways and conference bridges

.Figure 3 demonstrates LLDP-MED connectivity.

FIGURE 3 **LLDP-MED connectivity**

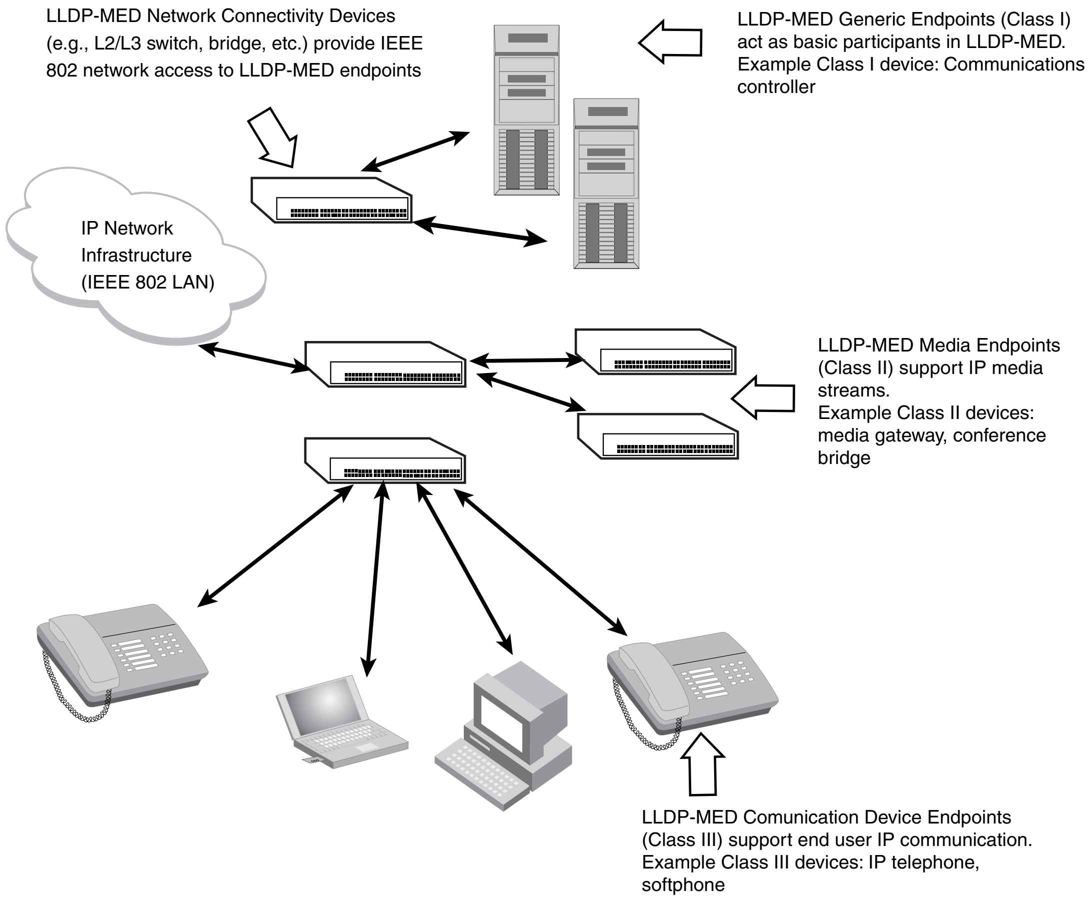

Benefits of LLDP-MED

LLDP-MED provides the following benefits:

- Vendor-independent management capabilities, enabling different IP telephony systems to interoperate in one network.
- Automatically deploys network policies, such as Layer 2 and Layer 3 QoS policies and Voice VLANs.
- Supports E-911 Emergency Call Services (ECS) for IP telephony
- Collects Endpoint inventory information
- Network troubleshooting
  - Helps to detect improper network policy configuration

LLDP-MED class

An LLDP-MED class specifies an Endpoint type and its capabilities. An Endpoint can belong to one of three LLDP-MED class types:

- Class 1 (Generic endpoint) – A Class 1 Endpoint requires basic LLDP discovery services, but does not support IP media nor does it act as an end-user communication appliance. A Class 1 Endpoint can be an IP communications controller, other communication-related server, or other device requiring basic LLDP discovery services.
- Class 2 (Media endpoint) – A Class 2 Endpoint supports media streams and may or may not be associated with a particular end user. Device capabilities include media streaming, as well as all of the capabilities defined for Class 1 Endpoints. A Class 2 Endpoint can be a voice/media gateway, conference, bridge, media server, etc.
- Class 3 (Communication endpoint) – A Class 3 Endpoint supports end user IP communication. Capabilities include aspects related to end user devices, as well as all of the capabilities defined for Class 1 and Class 2 Endpoints. A Class 3 Endpoint can be an IP telephone, softphone (PC-based phone), or other communication device that directly supports the end user.

  Discovery services defined in Class 3 include location identifier (ECS/E911) information and inventory management.

The LLDP-MED device class is advertised when LLDP-MED is enabled on a port.

Figure 3 illustrates LLDP-MED connectivity and supported LLDP-MED classes.

General LLDP operating principles

LLDP and LLDP-MED use the services of the Data Link sublayers, Logical Link Control and Media Access Control, to transmit and receive information to and from other *LLDP Agents* (protocol entities that implement LLDP).

LLDP is a one-way protocol. An LLDP agent can transmit and receive information to and from another LLDP agent located on an adjacent device, but it cannot solicit information from another LLDP agent, nor can it acknowledge information received from another LLDP agent.

LLDP operating modes

When LLDP is enabled on a global basis, by default, each port on the Brocade device will be capable of transmitting and receiving LLDP packets. You can disable a port's ability to transmit and receive LLDP packets, or change the operating mode to one of the following:

- Transmit LLDP information only
- Receive LLDP information only

LLDP transmit mode

An LLDP agent sends LLDP packets to adjacent LLDP-enabled devices. The LLDP packets contain information about the transmitting device and port.

An LLDP agent initiates the transmission of LLDP packets whenever the transmit countdown timing counter expires, or whenever LLDP information has changed. When a transmit cycle is initiated, the LLDP manager extracts the MIB objects and formats this information into TLVs. The TLVs are inserted into an LLDPDU, addressing parameters are prepended to the LLDPDU, and the information is sent out LLDP-enabled ports to adjacent LLDP-enabled devices.

LLDP receive mode

An LLDP agent receives LLDP packets from adjacent LLDP-enabled devices. The LLDP packets contain information about the transmitting device and port.

When an LLDP agent receives LLDP packets, it checks to ensure that the LLDPDUs contain the correct sequence of mandatory TLVs, then validates optional TLVs. If the LLDP agent detects any errors in the LLDPDUs and TLVs, it drops them in software. TLVs that are not recognized but do not contain basic formatting errors, are assumed to be valid and are assigned a temporary identification index and stored for future possible alter retrieval by network management. All validated TLVs are stored in the neighbor database.

LLDP packets

LLDP agents transmit information about a sending device/port in packets called LLDP Data Units (LLDPDUs). All the LLDP information to be communicated by a device is contained within a single 1500 byte packet. A device receiving LLDP packets is not permitted to combine information from multiple packets.

As shown in Figure 4, each LLDPDU has three mandatory TLVs, an End of LLDPDU TLV, plus optional TLVs as selected by network management.

FIGURE 4 **LLDPDU packet format**

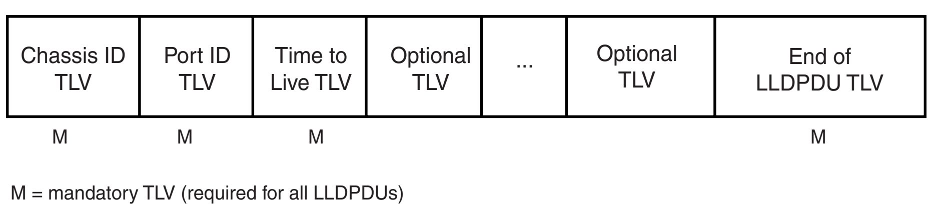

Each LLDPDU consists of an untagged Ethernet header and a sequence of short, variable length information elements known as type, length, value (TLV).

TLVs have Type, Length, and Value fields, where:

- **Type** identifies the kind of information being sent
- **Length** indicates the length (in octets) of the information string
- **Value** is the actual information being sent (for example, a binary bit map or an alpha-numeric string containing one or more fields).

TLV support

This section lists the LLDP and LLDP-MED TLV support.

LLDP TLVs

There are two types of LLDP TLVs, as specified in the IEEE 802.3AB standard:

- **Basic management TLVs** consist of both optional general system information TLVs as well as mandatory TLVs.

  Mandatory TLVs cannot be manually configured. They are always the first three TLVs in the LLDPDU, and are part of the packet header.

  General system information TLVs are optional in LLDP implementations and are defined by the Network Administrator.

  Brocade devices support the following Basic Management TLVs:

  - Chassis ID (mandatory)
  - Port ID (mandatory)
  - Time to Live (mandatory)
  - Port description
  - System name
  - System description
  - System capabilities
  - Management address
  - End of LLDPDU
- **Organizationally-specific TLVs** are optional in LLDP implementations and are defined and encoded by individual organizations or vendors. These TLVs include support for, but are not limited to, the IEEE 802.1 and 802.3 standards and the TIA-1057 standard.

  Brocade devices support the following Organizationally-specific TLVs:

  - **802.1 organizationally-specific TLVs**

    Port VLAN ID

    VLAN name TLV
  - **802.3 organizationally-specific TLVs**

    MAC/PHY configuration/status

    Power through MDI

    Link aggregation

    Maximum frame size

LLDP-MED TLVs

Brocade devices honor and send the following LLDP-MED TLVs, as defined in the TIA-1057 standard:

- LLDP-MED capabilities
- Network policy
- Location identification
- Extended power-via-MDI

Mandatory TLVs

When an LLDP agent transmits LLDP packets to other agents in the same 802 LAN segments, the following mandatory TLVs are always included:

- Chassis ID
- Port ID
- Time to Live (TTL)

This section describes the above TLVs in detail.

Chassis ID

The Chassis ID identifies the device that sent the LLDP packets.

There are several ways in which a device may be identified. A chassis ID subtype, included in the TLV and shown in Table 11, indicates how the device is being referenced in the Chassis ID field.

TABLE 11 Chassis ID subtypes

ID subtype	Description
0	Reserved
1	Chassis component
2	Interface alias
3	Port component
4	MAC address
5	Network address
6	Interface name
7	Locally assigned
8 – 255	Reserved

Brocade devices use chassis ID subtype 4, the base MAC address of the device. Other third party devices may use a chassis ID subtype other than 4. The chassis ID will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Chassis ID (MAC address):  0000.0033.e2c0
```

The chassis ID TLV is always the first TLV in the LLDPDU.

Port ID

The Port ID identifies the port from which LLDP packets were sent.

There are several ways in which a port may be identified, as shown in Figure 12. A port ID subtype, included in the TLV, indicates how the port is being referenced in the Port ID field.

TABLE 12 Port ID subtypes

ID subtype	Description
0	Reserved
1	Interface alias
2	Port component
3	MAC address
4	Network address
5	Interface name
6	Agent circuit ID
7	Locally assigned
8 – 255	Reserved

Brocade devices use port ID subtype 3, the permanent MAC address associated with the port. Other third party devices may use a port ID subtype other than 3. The port ID appears similar to the following on the remote device, and in the CLI display output on the Brocade device (show lldp local-info).

```
Port ID (MAC address):  0000.0033.e2d3
```

The LLDPDU format is shown in "LLDPDU packet format" on page 194.

The Port ID TLV format is shown below.

FIGURE 5 Port ID TLV packet format

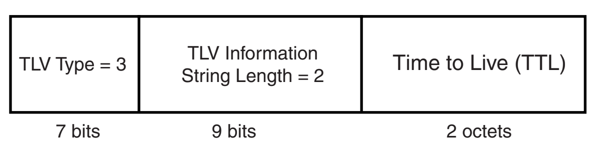

TTL value

The Time to Live (TTL) Value is the length of time the receiving device should maintain the information acquired by LLDP in its MIB.

The TTL value is automatically computed based on the LLDP configuration settings. The TTL value will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (show lldp local-info).

```
Time to live: 40 seconds
```

If the TTL field has a value other than zero, the receiving LLDP agent is notified to completely replace all information associated with the LLDP agent/port with the information in the received LLDPDU.

If the TTL field value is zero, the receiving LLDP agent is notified that all system information associated with the LLDP agent/port is to be deleted. This TLV may be used, for example, to signal that the sending port has initiated a port shutdown procedure.

The LLDPDU format is shown in "LLDPDU packet format" on page 194.

The TTL TLV format is shown below.

FIGURE 6 **TTL TLV packet format**

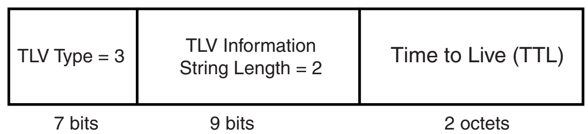

MIB support

Brocade devices support the following standard management information base (MIB) modules:

- LLDP-MIB
- LLDP-EXT-DOT1-MIB
- LLDP-EXT-DOT3-MIB
- LLDP-EXT-MED-MIB

Syslog messages

Syslog messages for LLDP provide management applications with information related to MIB data consistency and general status. These Syslog messages correspond to the lldpRemTablesChange SNMP notifications. Refer to "Enabling LLDP SNMP notifications and Syslog messages" on page 203.

Syslog messages for LLDP-MED provide management applications with information related to topology changes. These Syslog messages correspond to the lldpXMedTopologyChangeDetected SNMP notifications. Refer to "Enabling SNMP notifications and Syslog messages for LLDP-MED topology changes" on page 213.

LLDP configuration

This section describes how to enable and configure LLDP.

Table 13 lists the LLDP global-level tasks and the default behavior/value for each task.

TABLE 13 LLDP global configuration tasks and default behavior /value

Global task	Default behavior / value when LLDP is enabled
Enabling LLDP on a global basis	Disabled
Specifying the maximum number of LLDP neighbors per device	Automatically set to 392 neighbors per device
Specifying the maximum number of LLDP neighbors per port	Automatically set to 4 neighbors per port
Enabling SNMP notifications and Syslog messages	Disabled
Changing the minimum time between SNMP traps and Syslog messages	Automatically set to 2 seconds when SNMP notifications and Syslog messages for LLDP are enabled
Enabling and disabling TLV advertisements	When LLDP transmit is enabled, by default, the Brocade device will automatically advertise LLDP capabilities, except for the system description, VLAN name, and power-via-MDI information, which may be configured by the system administrator. Also, if desired, you can disable the advertisement of individual TLVs.
Changing the minimum time between LLDP transmissions	Automatically set to 2 seconds
Changing the interval between regular LLDP transmissions	Automatically set to 30 seconds
Changing the holdtime multiplier for transmit TTL	Automatically set to 4
Changing the minimum time between port reinitializations	Automatically set to 2 seconds

LLDP configuration notes and considerations

- LLDP is supported on Ethernet interfaces only.
- If a port is 802.1X-enabled, the transmission and reception of LLDP packets will only take place while the port is authorized.
- Cisco Discovery Protocol (CDP) and Brocade Discovery Protocol (FDP) run independently of LLDP. Therefore, these discovery protocols can run simultaneously on the same device.
- By default, the Brocade device limits the number of neighbors per port to four, and staggers the transmission of LLDP packets on different ports, in order to minimize any high-usage spikes to the CPU.
- By default, the Brocade device forwards
- Ports that are in blocking mode (spanning tree) can still receive LLDP packets from a forwarding port.
- Auto-negotiation status indicates what is being advertised by the port for 802.3 auto-negotiation.

Enabling and disabling LLDP

LLDP is enabled by default on individual ports. However, to run LLDP, you must first enable it on a global basis (on the entire device).

To enable LLDP globally, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)#lldp run
```

Syntax: [**no**] **lldp run**

Enabling support for tagged LLDP packets

By default, Brocade devices do not accept tagged LLDP packets from other vendors' devices. To enable support, apply the command **lldp tagged-packets process** at the Global CONFIG level of the CLI. When enabled, the device will accept incoming LLDP tagged packets if the VLAN tag matches any of the following:

- a configured VLAN on the port
- the default VLAN for a tagged port
- the configured untagged VLAN for a dual-mode port

To enable support for tagged LLDP packets, enter the following command.

```
Brocade(config)#lldp tagged-packets process
```

Syntax: [**no**] **lldp tagged-packets process**

Changing a port LLDP operating mode

When LLDP is enabled on a global basis, by default, each port on the Brocade device will be capable of transmitting and receiving LLDP packets. You can disable a port's ability to transmit and receive LLDP packets, or change the operating mode to one of the following:

- Transmit LLDP information only
- Receive LLDP information only

You can configure a different operating mode for each port on the Brocade device. For example, you could disable the receipt and transmission of LLDP packets on port e 2/1, configure port e 2/3 to only receive LLDP packets, and configure port e 2/5 to only transmit LLDP packets.

The following sections show how to change the operating mode.

Enabling and disabling receive and transmit mode

To disable the receipt and transmission of LLDP packets on individual ports, enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#no lldp enable ports e 2/4 e 2/5
```

The above command disables LLDP on ports 2/4 and 2/5. These ports will not transmit nor receive LLDP packets.

To enable LLDP on a port after it has been disabled, enter the following command.

```
Brocade(config)#lldp enable ports e 2/4
```

Syntax: [**no**] **lldp enable ports ethernet** *port-list* | **all**

Use the [no] form of the command to disable the receipt and transmission of LLDP packets on a port.

NOTE
When a port is configured to both receive and transmit LLDP packets and the MED capabilities TLV is enabled, LLDP-MED is enabled as well. LLDP-MED is not enabled if the operating mode is set to receive only or transmit only.

Enabling and disabling receive only mode

When LLDP is enabled on a global basis, by default, each port on the Brocade device will be capable of transmitting and receiving LLDP packets. To change the LLDP operating mode from receive and transmit mode to receive only mode, simply disable the transmit mode. Enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#no lldp enable transmit ports e 2/4 e 2/5 e 2/6
```

The above command changes the LLDP operating mode on ports 2/4, 2/5, and 2/6 from transmit and receive mode to receive only mode.

To change a port LLDP operating mode from transmit only to receive only, first disable the transmit only mode, then enable the receive only mode. Enter commands such as the following.

```
Brocade(config)#no lldp enable transmit ports e 2/7 e 2/8 e 2/9
Brocade(config)#lldp enable receive ports e 2/7 e 2/8 e 2/9
```

The above commands change the LLDP operating mode on ports 2/7, 2/8, and 2/9, from transmit only to receive only. Note that if you do not disable the transmit only mode, you will configure the port to both transmit and receive LLDP packets.

NOTE
LLDP-MED is not enabled when you enable the receive only operating mode. To enable LLDP-MED, you must configure the port to both receive and transmit LLDP packets. Refer to “Enabling and disabling receive and transmit mode” on page 200.

Syntax: [**no**] **lldp enable receive ports ethernet** *port-list* | **all**

Use the [no] form of the command to disable the receive only mode.

Enabling and Disabling Transmit Only Mode

When LLDP is enabled on a global basis, by default, each port on the Brocade device will be capable of transmitting and receiving LLDP packets. To change the LLDP operating mode to transmit only mode, simply disable the receive mode. Enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#no lldp enable receive ports e 2/4 e 2/5 e 2/6
```

The above command changes the LLDP operating mode on ports 2/4, 2/5, and 2/6 from transmit and receive mode to transmit only mode. Any incoming LLDP packets will be dropped in software.

To change a port LLDP operating mode from receive only to transmit only, first disable the receive only mode, then enable the transmit only mode. For example, enter commands such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#no lldp enable receive ports e 2/7 e 2/8
Brocade(config)#lldp enable transmit ports e 2/7 e 2/8
```

The above commands change the LLDP operating mode on ports 2/7 and 2/8 from receive only mode to transmit only mode. Any incoming LLDP packets will be dropped in software. Note that if you do not disable receive only mode, you will configure the port to both receive and transmit LLDP packets.

NOTE
LLDP-MED is not enabled when you enable the transmit only operating mode. To enable LLDP-MED, you must configure the port to both receive and transmit LLDP packets. Refer to "Enabling and disabling receive and transmit mode" on page 200.

Syntax: [**no**] **lldp enable transmit ports ethernet** *port-list* | **all**

Use the [no] form of the command to disable the *transmit only* mode.

Configuring LLDP processing on 802.1x blocked port

This feature adds support for reception and transmission of Link Layer Discovery Protocol (LLDP) packets over an 802.1x blocked port. The default behavior is to drop received LLDP packets and not to transmit LLDP packets over an 802.1x disabled port. To receive or transmit LLDP packets over 802.1x blocked port or in other words to enable the LLDP processing on 802.1x blocked ports, use the **lldp-pass-through** configuration command.

Enabling LLDP processing on 802.1x blocked port

To enable the LLDP processing on all 802.1x blocked ports, enter the following command at the 802.1X configuration mode:

```
Brocade(config-dot1x)# lldp-pass-through all
```

Syntax: [no] lldp-pass-through all

To enable LLDP processing on a specific 802.1x blocked port, enter the following command at the 802.1X configuration mode:

```
Brocade(config-dot1x)# lldp-pass-through ethernet 1/1/1
```

Syntax: [no] lldp-pass-through ethernet *port*

Specify the *port* variable in the format stackable switches-*stack-unit/slotnum/portnum*

The **no** form of these commands disables LLDP processing on 802.1x blocked ports.

For more information on LLDP and 801.1x, refer IEEE 802.1AB and IEEE 802.1x.

NOTE
If **lldp-pass-through** is disabled, the neighboring information is lost only after LLDP timeout period (default is 120).

Maximum number of LLDP neighbors

You can change the limit of the number of LLDP neighbors for which LLDP data will be retained, per device as well as per port.

Specifying the maximum number of LLDP neighbors per device

You can change the maximum number of neighbors for which LLDP data will be retained for the entire system.

For example, to change the maximum number of LLDP neighbors for the entire device to 26, enter the following command.

```
Brocade(config)#lldp max-total-neighbors 26
```

Syntax: [**no**] **lldp max-total-neighbors** *value*

Use the [no] form of the command to remove the static configuration and revert to the default value of 392.

where *value* is a number between 16 and 8192. The default number of LLDP neighbors per device is 392.

Use the **show lldp** command to view the configuration.

Specifying the maximum number of LLDP neighbors per port

You can change the maximum number of LLDP neighbors for which LLDP data will be retained for each port. By default, the maximum number is four and you can change this to a value between one and 64.

For example, to change the maximum number of LLDP neighbors to six, enter the following command.

```
Brocade(config)#lldp max-neighbors-per-port 6
```

Syntax: [**no**] **lldp max-neighbors-per-port** *value*

Use the [no] form of the command to remove the static configuration and revert to the default value of four.

where *value* is a number from 1 to 64. The default is number of LLDP neighbors per port is four.

Use the **show lldp** command to view the configuration.

Enabling LLDP SNMP notifications and Syslog messages

SNMP notifications and Syslog messages for LLDP provide management applications with information related to MIB data updates and general status.

When you enable LLDP SNMP notifications, corresponding Syslog messages are enabled as well. When you enable LLDP SNMP notifications, the device will send traps and corresponding Syslog messages whenever there are changes to the LLDP data received from neighboring devices.

LLDP SNMP notifications and corresponding Syslog messages are disabled by default. To enable them, enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#lldp enable snmp notifications ports e 4/2 to 4/6
```

The above command enables SNMP notifications and corresponding Syslog messages on ports 4/2 and 4/6. By default, the device will send no more than one SNMP notification and Syslog message within a five second period. If desired, you can change this interval. Refer to "Specifying the minimum time between SNMP traps and Syslog messages" on page 204.

Syntax: [**no**] **lldp enable snmp notifications ports ethernet** *port-list* | **all**

Specifying the minimum time between SNMP traps and Syslog messages

When SNMP notifications and Syslog messages for LLDP are enabled, the device will send no more than one SNMP notification and corresponding Syslog message within a five second period. If desired, you can throttle the amount of time between transmission of SNMP traps (lldpRemTablesChange) and Syslog messages from five seconds up to a value equal to one hour (3600 seconds).

NOTE
Because LLDP Syslog messages are rate limited, some LLDP information given by the system will not match the current LLDP statistics (as shown in the **show lldp statistics** command output).

To change the minimum time interval between traps and Syslog messages, enter a command such as the following.

```
Brocade(config)#lldp snmp-notification-interval 60
```

When the above command is applied, the LLDP agent will send no more than one SNMP notification and Syslog message every 60 seconds.

Syntax: [**no**] **lldp snmp-notification-interval** *seconds*

where *seconds* is a value between 5 and 3600. The default is 5 seconds.

Changing the minimum time between LLDP transmissions

The LLDP transmit delay timer limits the number of LLDP frames an LLDP agent can send within a specified time frame. When you enable LLDP, the system automatically sets the LLDP transmit delay timer to two seconds. If desired, you can change the default behavior from two seconds to a value between 1 and 8192 seconds.

NOTE
The LLDP transmit delay timer must not be greater than one quarter of the LLDP transmission interval (CLI command **lldp transmit-interval**).

The LLDP transmit delay timer prevents an LLDP agent from transmitting a series of successive LLDP frames during a short time period, when rapid changes occur in LLDP. It also increases the probability that multiple changes, rather than single changes, will be reported in each LLDP frame.

To change the LLDP transmit delay timer, enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#lldp transmit-delay 7
```

The above command causes the LLDP agent to wait a minimum of seven seconds after transmitting an LLDP frame and before sending another LLDP frame.

Syntax: [**no**] **lldp transmit-delay** *seconds*

where *seconds* is a value between 1 and 8192. The default is two seconds. Note that this value must not be greater than one quarter of the LLDP transmission interval (CLI command **lldp transmit-interval**).

Changing the interval between regular LLDP transmissions

The LLDP transmit interval specifies the number of seconds between regular LLDP packet transmissions. When you enable LLDP, by default, the device will wait 30 seconds between regular LLDP packet transmissions. If desired, you can change the default behavior from 30 seconds to a value between 5 and 32768 seconds.

To change the LLDP transmission interval, enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#lldp transmit-interval 40
```

The above command causes the LLDP agent to transmit LLDP frames every 40 seconds.

Syntax: [**no**] **lldp transmit-interval** *seconds*

where *seconds* is a value from 5 to 32768. The default is 30 seconds.

NOTE
Setting the transmit interval or transmit holdtime multiplier, or both, to inappropriate values can cause the LLDP agent to transmit LLDPDUs with TTL values that are excessively high. This in turn can affect how long a receiving device will retain the information if it is not refreshed.

Changing the holdtime multiplier for transmit TTL

The holdtime multiplier for transmit TTL is used to compute the actual time-to-live (TTL) value used in an LLDP frame. The TTL value is the length of time the receiving device should maintain the information in its MIB. When you enable LLDP, the device automatically sets the holdtime multiplier for TTL to four. If desired, you can change the default behavior from four to a value between two and ten.

To compute the TTL value, the system multiplies the LLDP transmit interval by the holdtime multiplier. For example, if the LLDP transmit interval is 30 and the holdtime multiplier for TTL is 4, then the value 120 is encoded in the TTL field in the LLDP header.

To change the holdtime multiplier, enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#lldp transmit-hold 6
```

Syntax: [**no**] **lldp transmit-hold** *value*

where *value* is a number from 2 to 10. The default value is 4.

NOTE
Setting the transmit interval or transmit holdtime multiplier, or both, to inappropriate values can cause the LLDP agent to transmit LLDPDUs with TTL values that are excessively high. This in turn can affect how long a receiving device will retain the information if it is not refreshed.

Changing the minimum time between port reinitializations

The LLDP re-initialization delay timer specifies the minimum number of seconds the device will wait from when LLDP is disabled on a port, until it will honor a request to re-enable LLDP on that port. When you enable LLDP, the system sets the re-initialization delay timer to two seconds. If desired, you can change the default behavior from two seconds to a value between one and ten seconds.

To set the re-initialization delay timer, enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#lldp reinit-delay 5
```

The above command causes the device to wait five seconds after LLDP is disabled, before attempting to honor a request to re-enable it.

Syntax: [**no**] **lldp reinit-delay** *seconds*

where *seconds* is a value from 1 – 10. The default is two seconds.

LLDP TLVs advertised by the Brocade device

When LLDP is enabled on a global basis, the Brocade device will automatically advertise the following information, except for the features noted:

General system information:

- Management address
- Port description
- System capabilities
- System description (not automatically advertised)
- System name

802.1 capabilities:

- VLAN name (not automatically advertised)
- Untagged VLAN ID

802.3 capabilities:

- Link aggregation information
- MAC/PHY configuration and status
- Maximum frame size
- Power-via-MDI information (not automatically advertised)

The above TLVs are described in detail in the following sections.

NOTE
The system description, VLAN name, and power-via-MDI information TLVs are not automatically enabled. The following sections show how to enable these advertisements.

General system information for LLDP

Except for the system description, the Brocade device will advertise the following system information when LLDP is enabled on a global basis:

- Management address
- Port description
- System capabilities
- System description (not automatically advertised)
- System name

Management Address

A management address is normally an IPv4 or IPv6 address that can be used to manage the device. Management address advertising has two modes: default, or explicitly configured. The default mode is used when no addresses are configured to be advertised for a given port. If any addresses are configured to be advertised for a given port, then only those addresses are advertised. This applies across address types, so for example, if just one IPv4 address is explicitly configured to be advertised for a port, then no IPv6 addresses will be advertised for that port (since none were configured to be advertised), even if IPv6 addresses are configured within the system.

If no management address is explicitly configured to be advertised, the Brocade device will use the first available IPv4 address and the first available IPv6 address (so it may advertise IPv4, IPv6 or both). A Layer 3 switch will select the first available address of each type from those configured on the following types of interfaces, in the following order of preference:

- Physical port on which LLDP will be transmitting the packet
- Virtual router interface (VE) on a VLAN that the port is a member of
- Dedicated management port
- Loop back interface
- Virtual router interface (VE) on any other VLAN
- Other physical port
- Other interface

For IPv6 addresses, link-local and anycast addresses will be excluded from these searches.

If no IP address is configured on any of the above, the port's current MAC address will be advertised.

To advertise a IPv4 management address, enter a command such as the following:

```
Brocade(config)#lldp advertise management-address ipv4 10.157.2.1 ports e 1/4
```

The management address will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**):

```
Management address (IPv4): 10.157.2.1
```

Syntax: [**no**] **lldp advertise management-address ipv4** *ipv4 address* **ports ethernet** *port list* | **all**

To support an IPv6 management address, there is a similar command that has equivalent behavior as the IPv4 command.

To advertise an IPv6 management address, enter a command such as the following:

```
Brocade(config)#lldp advertise management-address ipv6 2001:DB8::90 ports e 2/7
```

Syntax: [**no**] **lldp advertise management-address ipv6** *ipv6 address* **ports ethernet** *port list* | **all**

ipv4 address or *ipv6 address* or both are the addresses that may be used to reach higher layer entities to assist discovery by network management. In addition to management addresses, the advertisement will include the system interface number associated with the management address.

For *port list*, specify the port(s) in the format [*slotnum/*]*portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually; use the keyword to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Port description

The port description TLV identifies the port from which the LLDP agent transmitted the advertisement. The port description is taken from the ifDescr MIB object from MIB-II.

By default, the port description is automatically advertised when LLDP is enabled on a global basis. To disable advertisement of the port description, enter a command such as the following.

```
Brocade(config)#no lldp advertise port-description ports e 2/4 to 2/12
```

The port description will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Port description: "GigabitEthernet20"
```

Syntax: [**no**] **lldp advertise port-description ports ethernet** *port-list* | **all**

System capabilities

The system capabilities TLV identifies the primary functions of the device and indicates whether these primary functions are enabled. The primary functions can be one or more of the following (more than one for example, if the device is both a bridge and a router):

- Repeater
- Bridge
- WLAN access point
- Router
- Telephone
- DOCSIS cable device
- Station only (devices that implement end station capability)
- Other

System capabilities for Brocade devices are based on the type of software image in use (e.g., Layer 2 switch or Layer 3 router). The enabled capabilities will be the same as the available capabilities, except that when using a router image (base or full Layer 3), if the global route-only feature is turned on, the bridge capability will not be included, since no bridging takes place.

By default, the system capabilities are automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
Brocade(config)#no lldp advertise system-capabilities ports e 2/4 to 2/12
```

The system capabilities will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
System capabilities :   bridge
Enabled capabilities:   bridge
```

Syntax: [**no**] **lldp advertise system-capabilities ports ethernet** *port-list* | **all**

System description

The system description is the network entity, which can include information such as the product name or model number, the version of the system hardware type, the software operating system level, and the networking software version. The information corresponds to the sysDescr MIB object in MIB-II.

To advertise the system description, enter a command such as the following.

```
Brocade(config)#lldp advertise system-description ports e 2/4 to 2/12
```

The system description will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ System description  : "Brocade Communications,
Inc.,FCX_ADV_ROUTER_SOFT_PACKAGE,
IronWare Version 07.3.00T7f3 compiled on Sep 26 2011 at
       21:15:14 labeled as FCXR07300
```

NOTE
The contents of the show command output will vary depending on which TLVs are configured to be advertised.

Syntax: [**no**] **lldp advertise system-description ports ethernet** *port-list* | **all**

System name

The system name is the system administratively assigned name, taken from the sysName MIB object in MIB-II. The sysName MIB object corresponds to the name defined with the CLI command **hostname.**

By default, the system name is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
Brocade(config)#no lldp advertise system-name ports e 2/4 to 2/12
```

The system name will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
System name:  “FCX624SHPOE-ADV Router”
```

Syntax: [**no**] **lldp advertise system-name ports ethernet** *port-list* | **all**

802.1 capabilities

Except for the VLAN name, the Brocade device will advertise the following 802.1 attributes when LLDP is enabled on a global basis:

- VLAN name (not automatically advertised)

- Untagged VLAN ID

VLAN name

The VLAN name TLV contains the name and VLAN ID of a VLAN configured on a port. An LLDPDU may include multiple instances of this TLV, each for a different VLAN.

To advertise the VLAN name, enter a command such as the following.

```
Brocade(config)#lldp advertise vlan-name vlan 99 ports e 2/4 to 2/12
```

The VLAN name will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
VLAN name (VLAN 99): "Voice-VLAN-99"
```

Syntax: [**no**] **lldp advertise vlan-name vlan** *vlan ID* **ports ethernet** *port-list* | **all**

For *vlan ID*, enter the VLAN ID to advertise.

Untagged VLAN ID

The port VLAN ID TLV advertises the Port VLAN Identifier (PVID) that will be associated with untagged or priority-tagged frames. If the port is not an untagged member of any VLAN (i.e., the port is strictly a tagged port), the value zero will indicate that.

By default, the port VLAN ID is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
Brocade(config)#no lldp advertise port-vlan-id ports e 2/4 to 2/12
```

The untagged VLAN ID will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Port VLAN ID: 99
```

Syntax: [**no**] **lldp advertise port-vlan-id ports ethernet** *port-list* | **all**

802.3 capabilities

Except for Power-via-MDI information, the Brocade device will advertise the following 802.3 attributes when LLDP is enabled on a global basis:

- Link aggregation information
- MAC/PHY configuration and status
- Maximum frame size
- Power-via-MDI information (not automatically advertised)

Link aggregation TLV

The **link-aggregation** time, length, value (TLV) indicates the following:

- Whether the link is capable of being aggregated
- Whether the link is currently aggregated
- The primary trunk port

Brocade devices advertise link aggregation information about standard link aggregation (LACP) as well as static trunk configuration.

By default, link-aggregation information is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
Brocade(config)#no lldp advertise link-aggregation ports e 2/12
```

Syntax: [**no**] **lldp advertise link-aggregation ports ethernet** *port-list* | **all**

The link aggregation advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Link aggregation: not capable
```

MAC and PHY configuration status

The MAC and PHY configuration and status TLV includes the following information:

- Auto-negotiation capability and status
- Speed and duplex mode
- Flow control capabilities for auto-negotiation
- maximum port speed advertisement
- If applicable, indicates if the above settings are the result of auto-negotiation during link initiation or of a manual set override action

The advertisement reflects the effects of the following CLI commands:

- speed-duplex
- flow-control
- gig-default
- link-config

By default, the MAC/PHY configuration and status information are automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
Brocade(config)#no lldp advertise mac-phy-config-status ports e 2/4 to 2/12
```

The MAC/PHY configuration advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
 + 802.3 MAC/PHY       : auto-negotiation enabled
   Advertised capabilities: 10baseT-HD, 10baseT-FD, 100baseTX-HD,
100baseTX-FD,
   fdxSPause, fdxBPause, 1000baseT-HD, 1000baseT-FD
   Operational MAU type: 100BaseTX-FD
```

Syntax: [**no**] **lldp advertise mac-phy-config-status ports ethernet** *port-list* | **all**

Maximum frame size

The maximum frame size TLV provides the maximum 802.3 frame size capability of the port. This value is expressed in octets and includes the four-octet Frame Check Sequence (FCS). The default maximum frame size is 1522. The advertised value may change depending on whether the **aggregated-vlan** or **jumbo** CLI commands are in effect.

NOTE
On 48GC modules in non-jumbo mode, the maximum size of ping packets is 1486 bytes and the maximum frame size of tagged traffic is no larger than 1581 bytes.

By default, the maximum frame size is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
Brocade(config)#no lldp advertise max-frame-size ports e 2/4 to 2/12
```

The maximum frame size advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
Maximum frame size: 1522 octets
```

Syntax: [**no**] **lldp advertise max-frame-size ports ethernet** *port-list* | **all**

Power-via-MDI

The power-via-MDI TLV provides general information about Power over Ethernet (POE) capabilities and status of the port. It indicates the following:

- POE capability (supported or not supported)
- POE status (enabled or disabled)
- Power Sourcing Equipment (PSE) power pair – indicates which pair of wires is in use and whether the pair selection can be controlled. The Brocade implementation always uses pair A, and cannot be controlled.
- Power class – Indicates the range of power that the connected powered device has negotiated or requested.

NOTE
The power-via-MDI TLV described in this section applies to LLDP. There is also a power-via-MDI TLV for LLDP-MED devices, which provides extensive POE information. Refer to "Extended power-via-MDI information" on page 224.

To advertise the power-via-MDI information, enter a command such as the following.

```
Brocade(config)#lldp advertise power-via-mdi ports e 2/4 to 2/12
```

The power-via-MDI advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ 802.3 Power via MDI: PSE port, power enabled, class 0
  Power Pair         : A (not controllable)
```

Syntax: [**no**] **lldp advertise power-via-mdi ports ethernet** *port-list* | **all**

LLDP-MED configuration

This section provides the details for configuring LLDP-MED.

Table 14 lists the global and interface-level tasks and the default behavior/value for each task.

TABLE 14 LLDP-MED configuration tasks and default behavior / value

Task	Default behavior / value
Global CONFIG-level tasks	
Enabling LLDP-MED on a global basis	Disabled
Enabling SNMP notifications and Syslog messages for LLDP-MED topology change	Disabled
Changing the Fast Start Repeat Count	The system automatically sets the fast start repeat count to 3 when a Network Connectivity Device receives an LLDP packet from an Endpoint that is newly connected to the network. **NOTE:** The LLDP-MED fast start mechanism is only intended to run on links between Network Connectivity devices and Endpoint devices. It does not apply to links between LAN infrastructure elements, including between Network Connectivity devices, or to other types of links.
Interface-level tasks	
Defining a location ID	Not configured
Defining a network policy	Not configured

Enabling LLDP-MED

When LLDP is enabled globally, LLDP-MED is enabled if the LLDP-MED capabilities TLV is also enabled. By default, the LLDP-MED capabilities TLV is automatically enabled. To enable LLDP, refer to "Enabling and disabling LLDP" on page 200.

NOTE
LLDP-MED is not enabled on ports where the LLDP operating mode is receive only or transmit only. LLDP-MED is enabled on ports that are configured to both receive and transmit LLDP packets and have the LLDP-MED capabilities TLV enabled.

Enabling SNMP notifications and Syslog messages for LLDP-MED topology changes

SNMP notifications and Syslog messages for LLDP-MED provide management applications with information related to topology changes. For example, SNMP notifications can alert the system whenever a remote Endpoint device is connected to or removed from a local port. SNMP notifications identify the local port where the topology change occurred, as well as the device capability of the remote Endpoint device that was connected to or removed from the port.

When you enable LLDP-MED SNMP notifications, corresponding Syslog messages are enabled as well. When you enable LLDP-MED SNMP notifications, the device will send traps and Syslog messages when an LLDP-MED Endpoint neighbor entry is added or removed.

SNMP notifications and corresponding Syslog messages are disabled by default. To enable them, enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#lldp enable snmp med-topo-change-notifications ports e 4/4 to 4/6
```

Syntax: **no lldp enable snmp med-topo-change-notifications ports ethernet** *port-list* | **all**

Changing the fast start repeat count

The fast start feature enables a Network Connectivity Device to initially advertise itself at a faster rate for a limited time when an LLDP-MED Endpoint has been newly detected or connected to the network. This feature is important within a VoIP network, for example, where rapid availability is crucial for applications such as emergency call service location (E911).

The fast start timer starts when a Network Connectivity Device receives the first LLDP frame from a newly detected Endpoint.

The ***LLDP-MED fast start repeat count*** specifies the number of LLDP packets that will be sent during the LLDP-MED fast start period. By default, the device will send three packets at one-second intervals. If desired, you can change the number of packets the device will send per second, up to a maximum of 10.

NOTE
The LLDP-MED fast start mechanism is only intended to run on links between Network Connectivity devices and Endpoint devices. It does not apply to links between LAN infrastructure elements, including between Network Connectivity devices, or to other types of links.

To change the LLDP-MED fast start repeat count, enter commands such as the following.

```
Brocade(config)#lldp med fast-start-repeat-count 5
```

The above command causes the device to send five LLDP packets during the LLDP-MED fast start period.

Syntax: [**no**] **lldp med fast-start-repeat-count** *value*

where value is a number from 1 to 10, which specifies the number of packets that will be sent during the LLDP-MED fast start period. The default is 3.

Defining a location id

The LLDP-MED Location Identification extension enables the Brocade device to set the physical location that an attached Class III Endpoint will use for location-based applications. This feature is important for applications such as IP telephony, for example, where emergency responders need to quickly determine the physical location of a user in North America that has just dialed 911.

For each port, you can define one or more of the following location ID formats:

- Geographic location (coordinate-based)
- Civic address
- Emergency Call Services (ECS) Emergency Location Identification Number (ELIN)

The above location ID formats are defined in the following sections.

Coordinate-based location

Coordinate-based location is based on the IETF RFC 3825 [6] standard, which specifies a Dynamic Host Configuration Protocol (DHCP) option for the coordinate-based geographic location of a client.

When you configure an Endpoint location information using the coordinate-based location, you specify the latitude, longitude, and altitude, along with ***resolution indicators*** (a measure of the accuracy of the coordinates), and the reference ***datum*** (the map used for the given coordinates).

To configure a coordinate-based location for an Endpoint device, enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#lldp med location-id coordinate-based latitude
-78.303 resolution 20 longitude 34.27 resolution 18 altitude meters 50 resolution
16 wgs84
```

Syntax: [**no**] **lldp med location-id coordinate-based**
latitude *degrees* **resolution** *bits*
longitude *degrees* **resolution** *bits*
altitude floors *number* **resolution** *bits* | **meters** *number* **resolution** *bits*
datum

latitude *degrees* is the angular distance north or south from the earth equator measured through 90 degrees. Positive numbers indicate a location north of the equator and negative numbers indicate a location south of the equator.

resolution *bits* specifies the precision of the value given for latitude. A smaller value increases the area within which the device is located. For latitude, enter a number between 1 and 34.

longitude *degrees* is the angular distance from the intersection of the zero meridian. Positive values indicate a location east of the prime meridian and negative numbers indicate a location west of the prime meridian.

resolution *bits* specifies the precision of the value given for longitude. A smaller value increases the area within which the device is located. For longitude resolution, enter a number between 1 and 34.

altitude floors *number* is the vertical elevation of a building above the ground, where 0 represents the floor level associated with the ground level at the main entrance and larger values represent floors that are above (higher in altitude) floors with lower values. For example, 2 for the 2nd floor. Sub-floors can be represented by non-integer values. For example, a mezzanine between floor 1 and floor 2 could be represented as 1.1. Similarly, the mezzanines between floor 4 and floor 5 could be represented as 4.1 and 4.2 respectively. Floors located below ground level could be represented by negative values.

resolution *bits* specifies the precision of the value given for altitude. A smaller value increases the area within which the device is located. For floors resolution, enter the value 0 if the floor is unknown, or 30 if a valid floor is being specified.

altitude meters *number* is the vertical elevation in number of meters, as opposed to floors.

resolution *bits* specifies the precision of the value given for altitude. A smaller value increases the area within which the device is located. For meters resolution, enter a value from 0 to 30.

Datum is the map used as the basis for calculating the location. Specify one of the following:

- **wgs84** – (geographical 3D) – World Geodesic System 1984, CRS Code 4327, Prime Meridian Name: Greenwich
- **nad83-navd88** – North American Datum 1983, CRS Code 4269, Prime Meridian Name: Greenwich; The associated vertical datum is the North American Vertical Datum of 1988 (NAVD88). Use this datum when referencing locations on land. If land is near tidal water, use nad83-mllw (below).
- **nad83-mllw** – North American Datum 1983, CRS Code 4269, Prime Meridian Name: Greenwich; The associated vertical datum is mean lower low water (MLLW). Use this datum when referencing locations on water, sea, or ocean.

Example coordinate-based location configuration

The following shows an example coordinate-based location configuration for the Sears Tower, at the following location.

103rd Floor
233 South Wacker Drive
Chicago, IL 60606

```
Brocade(config)#lldp med location-id coordinate-based latitude 41.87884
resolution 18 longitude 87.63602 resolution 18 altitude floors 103 resolution 30
wgs84
```

The above configuration shows the following:

- Latitude is 41.87884 degrees north (or 41.87884 degrees).
- Longitude is 87.63602 degrees west (or 87.63602 degrees).
- The latitude and longitude resolution of 18 describes a geo-location area that is latitude 41.8769531 to latitude 41.8789062 and extends from -87.6367188 to -87.6347657 degrees longitude. This is an area of approximately 373412 square feet (713.3 ft. x 523.5 ft.).
- The location is inside a structure, on the 103rd floor.
- The WGS 84 map was used as the basis for calculating the location.

Example coordinate-based location advertisement

The coordinate-based location advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED Location ID
  Data Format: Coordinate-based
  Latitude Resolution  : 20 bits
  Latitude Value       : -78.303 degrees
  Longitude Resolution : 18 bits
  Longitude Value      : 34.27 degrees
  Altitude Resolution  : 16 bits
  Altitude Value       : 50. meters
  Datum                : WGS 84
```

Configuring civic address location

When you configure a media Endpoint location using the address-based location, you specify the location the entry refers to, the country code, and the elements that describe the civic or postal address.

To configure a civic address-based location for LLDP-MED, enter commands such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#lldp med location-id civic-address refers-to client country US
elem 1 CA elem 3 “Santa Clara” elem 6 “4980 Great America Pkwy” elem 24 95054 elem
27 5 elem 28 551 elem 29 office elem 23 “John Doe”
```

Syntax: [**no**] **lldp med location-id civic-address refers-to** *elem* **country** *country code* **elem** *CA type value* [**elem** *CA type value*] [**elem** *CA type value*]**....**

refers-to *elem* describes the location that the entry refers to. Specify one of the following:

- client
- dhcp-server
- network-element

where **dhcp-server** or **network-element** should only be used if it is known that the Endpoint is in close physical proximity to the DHCP server or network element.

country code is the two-letter ISO 3166 country code in capital ASCII letters.

Example

- CA – Canada
- DE – Germany
- JP – Japan
- KR – Korea
- US – United States

CA type is a value from 0 – 255, that describes the civic address element. For example, a CA type of 24 specifies a postal or zip code. Valid elements and their types are listed in Table 15.

value is the actual value of the elem *CA type*, above. For example, 95123 for the postal or zip code. Acceptable values are listed in Table 15, below.

NOTE
If the value of an element contains one or more spaces, use double quotation marks (“) at the beginning and end of the string. For example, `elem 3 "Santa Clara"`.

TABLE 15 Elements used with civic address

Civic Address (CA) type	Description	Acceptable values / examples
0	Language	The ISO 639 language code used for presenting the address information.
1	National subdivisions (state, canton, region, province, or prefecture)	Examples: Canada – Province Germany – State Japan – Metropolis Korea – Province United States – State
2	County, parish, gun (JP), or district (IN)	Examples: Canada – County Germany – County Japan – City or rural area Korea – County United States – County
3	City, township, or shi (JP)	Examples: Canada – City or town Germany – City Japan – Ward or village Korea – City or village United States – City or town

TABLE 15 Elements used with civic address (Continued)

Civic Address (CA) type	Description	Acceptable values / examples
4	City division, borough, city district, ward, or chou (JP)	Examples: Canada – N/A Germany – District Japan – Town Korea – Urban district United States – N/A
5	Neighborhood or block	Examples: Canada – N/A Germany – N/A Japan – City district Korea – Neighborhood United States – N/A
6	Street	Examples: Canada – Street Germany – Street Japan – Block Korea – Street United States – Street
16	Leading street direction	N (north), E (east), S (south), W (west), NE, NW, SE, SW
17	Trailing street suffix	N (north), E (east), S (south), W (west), NE, NW, SE, SW
18	Street suffix	Acceptable values for the United States are listed in the United States Postal Service Publication 28 [18], Appendix C. Example: Ave, Place
19	House number	The house number (street address) Example: 1234
20	House number suffix	A modifier to the house number. It does not include parts of the house number. Example: A, 1/2
21	Landmark or vanity address	A string name for a location. It conveys a common local designation of a structure, a group of buildings, or a place that helps to locate the place. Example: UC Berkeley
22	Additional location information	An unstructured string name that conveys additional information about the location. Example: west wing
23	Name (residence and office occupant)	Identifies the person or organization associated with the address. Example: Textures Beauty Salon
24	Postal / zip code	The valid postal / zip code for the address. Example: 95054-1234
25	Building (structure)	The name of a single building if the street address includes more than one building or if the building name is helpful in identifying the location. Example: Law Library

TABLE 15 Elements used with civic address (Continued)

Civic Address (CA) type	Description	Acceptable values / examples
26	Unit (apartment, suite)	The name or number of a part of a structure where there are separate administrative units, owners, or tenants, such as separate companies or families who occupy that structure. Common examples include suite or apartment designations. Example: Apt 27
27	Floor	Example: 4
28	Room number	The smallest identifiable subdivision of a structure. Example: 7A
29	Placetype	The type of place described by the civic coordinates. For example, a home, office, street, or other public space. Example: Office
30	Postal community name	When the postal community name is defined, the civic community name (typically CA type 3) is replaced by this value. Example: Alviso
31	Post office box (P.O. box)	When a P.O. box is defined, the street address components (CA types 6, 16, 17, 18, 19, and 20) are replaced with this value. Example: P.O. Box 1234
32	Additional code	An additional country-specific code that identifies the location. For example, for Japan, this is the Japan Industry Standard (JIS) address code. The JIS address code provides a unique address inside of Japan, down to the level of indicating the floor of the building.
128	Script	The script (from ISO 15924 [14]) used to present the address information. Example: Latn **NOTE:** If not manually configured, the system assigns the default value **Latn**
255	Reserved	

Example civic address location advertisement

The Civic address location advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED Location ID
  Data Format: Civic Address
  Location of: Client
  Country    : "US"
  CA Type    : 1
  CA Value   : "CA"
  CA Type    : 3
  CA Value   : "Santa Clara"
  CA Type    : 6
  CA Value   : "4980 Great America Pkwy."
  CA Type    : 24
  CA Value   : "95054"
  CA Type    : 27
  CA Value   : "5"
  CA Type    : 28
  CA Value   : "551"
  CA Type    : 29
  CA Value   : "office"
  CA Type    : 23
  CA Value   : "John Doe"
```

Configuring emergency call service

The Emergency Call Service (ECS) location is used specifically for Emergency Call Services applications.

When you configure a media Endpoint location using the emergency call services location, you specify the Emergency Location Identification Number (ELIN) from the North America Numbering Plan format, supplied to the Public Safety Answering Point (PSAP) for ECS purposes.

To configure an ECS-based location for LLDP-MED, enter a command such as the following at the Global CONFIG level of the CLI.

```
Brocade(config)#lldp med location-id ecs-elin 4082071700
```

Syntax: [**no**] **lldp med location-id ecs-elin** *number* **ports ethernet** *port-list* | **all**

number is a number from 10 to 25 digits in length.

Example ECS ELIN location advertisements

The ECS ELIN location advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED Location ID
  Data Format: ECS ELIN
  Value      : 4082071700
```

Defining an LLDP-MED network policy

An LLDP-MED network policy defines an Endpoint VLAN configuration (VLAN type and VLAN ID) and associated Layer 2 and Layer 3 priorities that apply to a specific set of applications on a port.

NOTE
This feature applies to applications that have specific real-time network policy requirements, such as interactive voice or video services. It is not intended to run on links other than between Network Connectivity devices and Endpoints, and therefore does not advertise the multitude of network policies that frequently run on an aggregated link.

To define an LLDP-MED network policy for an Endpoint, enter a command such as the following.

```
Brocade(config)#lldp med network-policy application voice tagged vlan 99 priority
3 dscp 22 port e 2/6
```

The network policy advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED Network Policy
 Application Type  : Voice
 Policy Flags      : Known Policy, Tagged
 VLAN ID           : 99
 L2 Priority       : 3
 DSCP Value        : 22
```

NOTE
Endpoints will advertise a policy as “unknown” in the **show lldp neighbor detail** command output, if it is a policy that is required by the Endpoint and the Endpoint has not yet received it.

LLDP-MED network policy configuration syntax

The CLI syntax for defining an LLDP-MED network policy differs for tagged, untagged, and priority tagged traffic. Refer to the appropriate syntax, below.

For tagged traffic

Syntax: [**no**] **lldp med network-policy application** *application type* **tagged vlan** *vlan ID* **priority** *0 – 7* **dscp** *0 – 63* **ports ethernet** *port-list* | **all**

For untagged traffic

Syntax: [**no**] **lldp med network-policy application** *application type* **untagged dscp** *0 – 63* **ports ethernet** *port-list* | **all**

For priority-tagged traffic

Syntax: [**no**] **lldp med network-policy application** *application type* **priority-tagged priority** *0 – 7* **dscp** *0 – 63* **ports ethernet** *port-list* | **all**

application type indicates the primary function of the applications defined by this network policy. Application type can be one of the following:

- **guest-voice** – Limited voice service for guest users and visitors with their own IP telephony handsets or similar devices that support interactive voice services.
- **guest-voice-signaling** – Limited voice service for use in network topologies that require a different policy for guest voice signaling than for guest voice media.
- **softphone-voice** – Softphone voice service for use with multi-media applications that work in association with VoIP technology, enabling phone calls direct from a PC or laptop. Softphones do not usually support multiple VLANs, and are typically configured to use an untagged VLAN or a single tagged data-specific VLAN. Note that when a network policy is defined for use with an untagged VLAN, the Layer 2 priority field is ignored and only the DSCP value is relevant.
- **streaming-video** – Applies to broadcast- or multicast-based video content distribution and similar applications that support streaming video services requiring specific network policy treatment. Video applications that rely on TCP without buffering would not be an intended use of this application type.
- **video-conferencing** – Applies to dedicated video conferencing equipment and similar devices that support real-time interactive video/audio services.
- **video-signaling** – For use in network topologies that require a separate policy for video signaling than for video media. Note that this application type should not be advertised if all the same network policies apply as those advertised in the video conferencing policy TLV.
- **voice** – For use by dedicated IP telephony handsets and similar devices that support interactive voice services.
- **voice-signaling** – For use in network topologies that require a different policy for voice signaling than for voice media. Note that this application type should not be advertised if all the same network policies apply as those advertised in the voice policy TLV.
- **tagged vlan** *vlan id* specifies the tagged VLAN that the specified application type will use.
- **untagged** indicates that the device is using an untagged frame format.
- **priority-tagged** indicates that the device uses priority-tagged frames. In this case, the device uses the default VLAN (PVID) of the ingress port.
- **priority** *0 –7* indicates the Layer 2 priority value to be used for the specified application type. Enter 0 to use the default priority.

- **dscp** *0 – 63* specifies the Layer 3 Differentiated Service codepoint priority value to be used for the specified application type. Enter 0 to use the default priority.

LLDP-MED attributes advertised by the Brocade device

LLDP-MED attributes are only advertised on a port if LLDP-MED is enabled (which is done by enabling the LLDP-MED capabilities TLV), the port operating mode is *receive* and *transmit* (the default), and the port has received an LLDP-MED advertisement from an Endpoint. By default, the Brocade device will automatically advertise the following LLDP-MED attributes when the above criteria are met:

- LLDP-MED capabilities
- Location ID
- Network policy
- Power-via-MDI information

NOTE
Although the Location ID and Network policy attributes are automatically advertised, they will have no effect until they are actually defined.

LLDP-MED capabilities

When enabled, LLDP-MED is enabled, and the LLDP-MED capabilities TLV is sent whenever any other LLDP-MED TLV is sent. When disabled, LLDP-MED is disabled and no LLDP-MED TLVs are sent.

The LLDP-MED capabilities advertisement includes the following information:

- The supported LLDP-MED TLVs
- The device type (Network Connectivity device or Endpoint (Class 1, 2, or 3))

By default, LLDP-MED information is automatically advertised when LLDP-MED is enabled. To disable this advertisement, enter a command such as the following.

```
Brocade(config)#no lldp advertise med-capabilities ports e 2/4 to 2/12
```

NOTE
Disabling the LLDP-MED capabilities TLV disables LLDP-MED.

To re-enable the LLDP-MED Capabilities TLV (and LLDP-MED) after it has been disabled, enter a command such as the following.

```
Brocade(config)#lldp advertise med-capabilities ports e 2/4 to 2/12
```

The LLDP-MED capabilities advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
  + MED capabilities: capabilities, networkPolicy, location, extendedPSE
MED device type : Network Connectivity
```

Syntax: [**no**] **lldp advertise med-capabilities ports ethernet** *port-list* | **all**

Extended power-via-MDI information

The extended Power-via-MDI TLV enables advanced power management between LLDP-MED Endpoints and Network Connectivity Devices. This TLV provides significantly more information than the 802.1AB Power-via-MDI TLV referenced in "Power-via-MDI" on page 212. For example, this TLV enables an Endpoint to communicate a more precise required power level, thereby enabling the device to allocate less power to the Endpoint, while making more power available to other ports.

The LLDP-MED Power-via-MDI TLV advertises an Endpoint IEEE 802.3af power-related information, including the following:

- Power type – indicates whether the LLDP-MED device transmitting the LLPDU is a power sourcing device or a powered device:
  - Power sourcing device/equipment (PSE) – This is the source of the power, or the device that integrates the power onto the network. Power sourcing devices/equipment have embedded POE technology. In this case, the power sourcing device is the Brocade POE device.
  - Powered device (PD) – This is the Ethernet device that requires power and is situated on the other end of the cable opposite the power sourcing device.
- Power source – The power source being utilized by a PSE or PD, for example, primary power source, backup power source, or unknown.

  For Endpoint devices, the power source information indicates the power capability of the Network Connectivity Device it is attached to. When the Network Connectivity device advertises that it is using its primary power source, the Endpoint should expect to have uninterrupted access to its available power. Likewise, if the Network Connectivity device advertises that it is using backup power, the Endpoint should not expect continuous power. The Endpoint may additionally choose to power down non-essential subsystems or to conserve power as long as the PSE is advertising that it is operating on backup power.

  NOTE
  Brocade devices always advertise the power source as "unknown".

- Power priority – The in-line power priority level for the PSE or PD:
  - 3 – low
  - 2 – high
  - 1 – critical
  - unknown
- Power level – The total power, in tenths of watts, required by a PD from a PSE, or the total power a PSE is capable of sourcing over a maximum length cable based on its current configuration.

  If the exact power is not known for a PSE or PD, it will advertise the power level associated with its 802.3af power class (listed in Table 16).

TABLE 16 802.3af power classes

Power class	Minimum power level output at the PSE	Maximum power levels at the PD
0	15.4 watts	0.44 – 12.95 watts
1	4.0 watts	0.44 – 3.84 watts

TABLE 16 802.3af power classes

Power class	Minimum power level output at the PSE	Maximum power levels at the PD
2	7.0 watts	3.84 – 6.49 watts
3	15.4 watts	6.49 – 12.95 watts

For a PD (Endpoint device), the power level represents the maximum power it can consume during normal operations in its current configuration, even if its actual power draw at that instance is less than the advertised power draw.

For a PSE (Network Connectivity device), the power level represents the amount of power that is available on the port at the time. If the PSE is operating in reduced power (i.e., it is using backup power), the reduced power capacity is advertised as long as the condition persists.

By default, LLDP-MED power-via-MDI information is automatically advertised when LLDP-MED is enabled, the port is a POE port, and POE is enabled on the port. To disable this advertisement, enter a command such as the following.

```
Brocade(config)#no lldp advertise med-power-via-mdi ports e 2/4 to 2/12
```

The LLDP-MED power-via-MDI advertisement will appear similar to the following on the remote device, and in the CLI display output on the Brocade device (**show lldp local-info**).

```
+ MED Extended Power via MDI
   Power Type     : PSE device
   Power Source   : Unknown Power Source
   Power Priority : Low (3)
   Power Value    : 6.5 watts (PSE equivalent: 7005 mWatts)
```

Syntax: [**no**] **lldp advertise med-power-via-mdi ports ethernet** *port-list* | **all**

Displaying LLDP statistics and configuration settings

You can use the following CLI **show** commands to display information about LLDP settings and statistics:

- **show lldp** – Displays a summary of the LLDP configuration settings.
- **show lldp statistics** – Displays LLDP global and per-port statistics.
- **show lldp neighbors** – Displays a list of the current LLDP neighbors.
- **show lldp neighbors detail** – Displays the details of the latest advertisements received from LLDP neighbors.
- **show lldp local-info** – Displays the details of the LLDP advertisements that will be transmitted on each port.

This above **show** commands are described in this section.

LLDP configuration summary

To display a summary of the LLDP configuration settings on the device, enter the **show lldp** command at any level of the CLI.

The following shows an example report.

```
Brocade#show lldp
LLDP transmit interval            : 10 seconds
LLDP transmit hold multiplier     : 4  (transmit TTL: 40 seconds)
LLDP transmit delay               : 1 seconds
LLDP SNMP notification interval   : 5 seconds
LLDP reinitialize delay           : 1 seconds
LLDP-MED fast start repeat count  : 3

LLDP maximum neighbors            : 392
LLDP maximum neighbors per port   : 4
```

Syntax: **show lldp**

The following table describes the information displayed by the **show lldp statistics** command.

Field	Description
LLDP transmit interval	The number of seconds between regular LLDP packet transmissions.
LLDP transmit hold multiplier	The multiplier used to compute the actual time-to-live (TTL) value of an LLDP advertisement. The TTL value is the transmit interval multiplied by the transmit hold multiplier.
LLDP transmit delay	The number of seconds the LLDP agent will wait after transmitting an LLDP frame and before transmitting another LLDP frame.
LLDP SNMP notification interval	The number of seconds between transmission of SNMP LLDP traps (lldpRemTablesChange) and SNMP LLDP-MED traps (lldpXMedTopologyChangeDetected).
LLDP reinitialize delay	The minimum number of seconds the device will wait from when LLDP is disabled on a port, until a request to re-enable LLDP on that port will be honored.
LLDP-MED fast start repeat count	The number of seconds between LLDP frame transmissions when an LLDP-MED Endpoint is newly detected.
LLDP maximum neighbors	The maximum number of LLDP neighbors for which LLDP data will be retained, per device.
LLDP maximum neighbors per port	The maximum number of LLDP neighbors for which LLDP data will be retained, per port.

Displaying LLDP statistics

The **show lldp statistics** command displays an overview of LLDP neighbor detection on the device, as well as packet counters and protocol statistics. The statistics are displayed on a global basis.

The following shows an example report.

```
Brocade#show lldp statistics
Last neighbor change time: 23 hours 50 minutes 40 seconds ago

Neighbor entries added          : 14
Neighbor entries deleted        : 5
Neighbor entries aged out       : 4
Neighbor advertisements dropped : 0

Port        Tx Pkts    Rx Pkts    Rx Pkts    Rx Pkts    Rx TLVs    Rx TLVs Neighbors
              Total      Total   w/Errors  Discarded Unrecognz  Discarded  Aged Out
1             60963      75179          0          0          0          0         4
2                 0          0          0          0          0          0         0
3             60963      60963          0          0          0          0         0
4             60963     121925          0          0          0          0         0
5                 0          0          0          0          0          0         0
6                 0          0          0          0          0          0         0
7                 0          0          0          0          0          0         0
8                 0          0          0          0          0          0         0
9                 0          0          0          0          0          0         0
10            60974          0          0          0          0          0         0
11                0          0          0          0          0          0         0
12                0          0          0          0          0          0         0
13                0          0          0          0          0          0         0
14                0          0          0          0          0          0         0
```

Syntax: **show lldp statistics**

NOTE
You can reset LLDP statistics using the CLI command **clear LLDP statistics**. Refer to "Resetting LLDP statistics" on page 231.

The following table describes the information displayed by the **show lldp statistics** command.

Field	Description
Last neighbor change time	The elapsed time (in hours, minutes, and seconds) since a neighbor last advertised information. For example, the elapsed time since a neighbor was last added, deleted, or its advertised information changed.
Neighbor entries added	The number of new LLDP neighbors detected since the last reboot or since the last time the **clear lldp statistics** all command was issued.
Neighbor entries deleted	The number of LLDP neighbors deleted since the last reboot or since the last time the **clear lldp statistics all** command was issued.
Neighbor entries aged out	The number of LLDP neighbors dropped on all ports after the time-to-live expired. Note that LLDP entries age out naturally when a port cable or module is disconnected or when a port becomes disabled. However, if a disabled port is re-enabled, the system will delete the old LLDP entries.
Neighbor advertisements dropped	The number of valid LLDP neighbors the device detected, but could not add. This can occur, for example, when a new neighbor is detected and the device is already supporting the maximum number of neighbors possible. This can also occur when an LLDPDU is missing a mandatory TLV or is not formatted correctly.
Port	The local port number.
Tx Pkts Total	The number of LLDP packets the port transmitted.
Rx Pkts Total	The number of LLDP packets the port received.

Field	Description
Rx Pkts w/Errors	The number of LLDP packets the port received that have one or more detectable errors.
Rx Pkts Discarded	The number of LLDP packets the port received then discarded.
Rx TLVs Unrecognz	The number of TLVs the port received that were not recognized by the LLDP local agent. Unrecognized TLVs are retained by the system and can be viewed in the output of the **show LLDP neighbors** detail command or retrieved through SNMP.
Rx TLVs Discarded	The number of TLVs the port received then discarded.
Neighbors Aged Out	The number of times a neighbor information was deleted because its TTL timer expired.

Displaying LLDP neighbors

The **show lldp neighbors** command displays a list of the current LLDP neighbors per port.

The following shows an example report.

```
Brocade#show lldp neighbors
Lcl Port Chassis ID      Port ID         Port Description        System Name
1         0000.0034.0fc0  0000.0034.0fc0  GigabitEthernet9/1      FastIron Supe~
1         0000.0001.4000  0000.0001.4000  GigabitEthernet0/1/1    FastIron SX Swi~
3         0000.0011.0200  0000.0011.0203  GigabitEthernet4        FastIron SX 8~
4         0000.0011.0200  0000.0011.0202  GigabitEthernet3        FastIron SX 8~
4         0000.0011.0200  0000.0011.0210  GigabitEthernet17       FastIron SX 8~
15        0000.0011.0200  0000.0011.020f  GigabitEthernet16       FastIron SX 8~
16        0000.0011.0200  0000.0011.020e  GigabitEthernet15       FastIron SX 8~
17        0000.0011.0200  0000.0011.0211  GigabitEthernet18       FastIron SX 8~
18        0000.0011.0200  0000.0011.0210  GigabitEthernet17       FastIron SX 8~
```

Syntax: **show lldp neighbors**

The following table describes the information displayed by the **show lldp neighbors** command.

Field	Description
Lcl Port	The local LLDP port number.
Chassis ID	The identifier for the chassis. Brocade devices use the base MAC address of the device as the Chassis ID.
Port ID	The identifier for the port. Brocade devices use the permanent MAC address associated with the port as the port ID.
Port Description	The description for the port. Brocade devices use the ifDescr MIB object from MIB-II as the port description.
System Name	The administratively-assigned name for the system. Brocade devices use the sysName MIB object from MIB-II, which corresponds to the CLI **hostname** command setting. **NOTE:** A tilde (~) at the end of a line indicates that the value in the field is too long to display in full and is truncated.

Displaying LLDP neighbors detail

The **show lldp neighbors detail** command displays the LLDP advertisements received from LLDP neighbors.

The following shows an example **show lldp neighbors detail** report.

NOTE
The **show lldp neighbors detail** output will vary depending on the data received. Also, values that are not recognized or do not have a recognizable format, may be displayed in hexadecimal binary form.

```
Brocade#show lldp neighbors detail ports e 1/9
Local port: 1/9
  Neighbor: 0000.0018.cc03, TTL 101 seconds
    + Chassis ID (network address): 10.43.39.151
    + Port ID (MAC address): 0000.0018.cc03
    + Time to live: 120 seconds
    + Port description    : "LAN port"
    + System name         : "regDN 1015,MITEL 5235 DM"
    + System description  : "regDN 1015,MITEL 5235 DM,h/w rev 2,ASIC rev 1,f/w\
                              Boot 02.01.00.11,f/w Main 02.01.00.11"
    + System capabilities : bridge, telephone
      Enabled capabilities: bridge, telephone
    + Management address (IPv4): 10.43.39.151
    + 802.3 MAC/PHY          : auto-negotiation enabled
      Advertised capabilities: 10BaseT-HD, 10BaseT-FD, 100BaseTX-HD,
                               100BaseTX-FD
      Operational MAU type   : 100BaseTX-FD
    + MED capabilities: capabilities, networkPolicy, extendedPD
      MED device type : Endpoint Class III
    + MED Network Policy
      Application Type  : Voice
      Policy Flags      : Known Policy, Tagged
      VLAN ID           : 300
      L2 Priority       : 7
      DSCP Value        : 7
    + MED Extended Power via MDI
      Power Type     : PD device
      Power Source   : Unknown Power Source
      Power Priority : High (2)
      Power Value    : 6.2 watts (PSE equivalent: 6656 mWatts)
    + MED Hardware revision : "PCB Version: 2"
    + MED Firmware revision : "Boot 02.01.00.11"
    + MED Software revision : "Main 02.01.00.11"
    + MED Serial number     : ""
    + MED Manufacturer      : "Mitel Corporation"
    + MED Model name        : "MITEL 5235 DM"
    + MED Asset ID          : ""
```

A backslash (\) at the end of a line indicates that the text continues on the next line.

Except for the following field, the fields in the above output are described in the individual TLV advertisement sections in this chapter.

Field	Description
Neighbor	The source MAC address from which the packet was received, and the remaining TTL for the neighbor entry.

Syntax: show lldp neighbors detail [**ports ethernet** *port-list* | **all**]

If you do not specify any ports or use the keyword **all**, by default, the report will show the LLDP neighbor details for all ports.

Displaying LLDP configuration details

The **show lldp local-info** command displays the local information advertisements (TLVs) that will be transmitted by the LLDP agent.

NOTE
The **show lldp local-info** output will vary based on LLDP configuration settings.

The following shows an example report.

```
Brocade#show lldp local-info ports e 20
Local port: 20
  + Chassis ID (MAC address): 0000.0033.e2c0
  + Port ID (MAC address): 0000.0033.e2d3
  + Time to live: 40 seconds
  + System name: "FCX624SHPOE-ADV Router"
  + Port description: "GigabitEthernet20"
  + System description  : "Brocade Communications, Inc.
   FCX_ADV_ROUTER_SOFT_PACKAGE,
   IronWare Version 07.3.00T7f3 compiled on Sep 26 2011 at 21:15:14 labeled as
                            FCXR07300"
  + System capabilities : bridge
    Enabled capabilities: bridge
  + 802.3 MAC/PHY       : auto-negotiation enabled
    Advertised capabilities: 10BaseT-HD, 10BaseT-FD, 100BaseTX-HD,
                             100BaseTX-FD, fdxSPause, fdxBPause, 1000BaseT-HD,
                             1000BaseT-FD
    Operational MAU type: 100BaseTX-FD
  + 802.3 Power via MDI: PSE port, power enabled, class 2
    Power Pair         : A (not controllable)
  + Link aggregation: not capable
  + Maximum frame size: 1522 octets
  + MED capabilities: capabilities, networkPolicy, location, extendedPSE
    MED device type : Network Connectivity
  + MED Network Policy
    Application Type  : Voice
    Policy Flags      : Known Policy, Tagged
    VLAN ID           : 99
    L2 Priority       : 3
    DSCP Value        : 22
  + MED Network Policy
    Application Type  : Video Conferencing
    Policy Flags      : Known Policy, Tagged
    VLAN ID           : 100
    L2 Priority       : 5
    DSCP Value        : 10
  + MED Location ID
    Data Format: Coordinate-based location
```

```
    Latitude Resolution  : 20 bits
    Latitude Value       : -78.303 degrees
    Longitude Resolution : 18 bits
    Longitude Value      : 34.27 degrees
    Altitude Resolution  : 16 bits
    Altitude Value       : 50. meters
    Datum                : WGS 84
+ MED Location ID
    Data Format: Civic Address
    Location of: Client
    Country    : "US"
    CA Type    : 1
    CA Value   : "CA"
    CA Type    : 3
    CA Value   : "Santa Clara"
    CA Type    : 6
    CA Value   : "4980 Great America Pkwy."
    CA Type    : 24
    CA Value   : "95054"
    CA Type    : 27
    CA Value   : "5"
    CA Type    : 28
    CA Value   : "551"
    CA Type    : 29
    CA Value   : "office"
    CA Type    : 23
    CA Value   : "John Doe"
  + MED Location ID
    Data Format: ECS ELIN
    Value      : "1234567890"
  + MED Extended Power via MDI
    Power Type     : PSE device
    Power Source   : Unknown Power Source
    Power Priority : Low (3)
    Power Value    : 6.5 watts (PSE equivalent: 7005 mWatts) + Port VLAN ID: 99
  + Management address (IPv4): 10.1.1.121
  + VLAN name (VLAN 99): "Voice-VLAN-99"
```

NOTE
The contents of the **show** output will vary depending on which TLVs are configured to be advertised.

A backslash (\) at the end of a line indicates that the text continues on the next line.

The fields in the above output are described in the individual TLV advertisement sections in this chapter.

Syntax: **show lldp local-info** [**ports ethernet** *port-list* | **all**]

If you do not specify any ports or use the keyword **all**, by default, the report will show the local information advertisements for all ports.

Resetting LLDP statistics

To reset LLDP statistics, enter the **clear lldp statistics** command at the Global CONFIG level of the CLI. The Brocade device will clear the global and per-port LLDP neighbor statistics on the device (refer to "Displaying LLDP statistics" on page 226).

```
Brocade#clear lldp statistics
```

Syntax: clear lldp statistics [**ports ethernet** *port-list* | **all**]

If you do not specify any ports or use the keyword **all**, by default, the system will clear lldp statistics on all ports.

Clearing cached LLDP neighbor information

The Brocade device clears cached LLDP neighbor information after a port becomes disabled and the LLDP neighbor information ages out. However, if a port is disabled then re-enabled before the neighbor information ages out, the device will clear the cached LLDP neighbor information when the port is re-enabled.

If desired, you can manually clear the cache. For example, to clear the cached LLDP neighbor information for port e 20, enter the following command at the Global CONFIG level of the CLI.

```
Brocade#clear lldp neighbors ports e 20
```

Syntax: clear lldp neighbors [**ports ethernet** *port-list* | **all**]

If you do not specify any ports or use the keyword **all**, by default, the system will clear the cached LLDP neighbor information for all ports.

Chapter

Hardware Component Monitoring 9

Table 17 lists the individual Brocade FastIron switches and the hardware monitoring features they support. These features are supported in the Layer 2 and Layer 3 software images.

TABLE 17 Supported hardware monitoring features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6430	ICX 6450
Virtual cable testing (VCT)	Yes	Yes	No	No	No
Digital optical monitoring	Yes	Yes	Yes	Yes	Yes

This chapter contains the following sections:

NOTE
VCT is not supported on SX-FI48GPP, SX-FI-24GPP, SX-FI-24HF, SX-FI-2XG, and SX-FI-8XG.

The procedures in this chapter describe how to configure the software to monitor hardware components.

Virtual cable testing

FastIron devices support **Virtual Cable Test** (VCT) technology. VCT technology enables the diagnosis of a conductor (wire or cable) by sending a pulsed signal into the conductor, then examining the reflection of that pulse. This method of cable analysis is referred to as Time Domain Reflectometry (TDR). By examining the reflection, the Brocade device can detect and report cable statistics such as local and remote link pair, cable length, and link status.

Virtual cable testing configuration notes

- This feature is supported on copper ports only. It is not supported on fiber ports.
- This feature is not supported on the SX-FI48GPP module running software release 07.2.02 or later.
- This feature is not supported on SX-FI2XG, SX-FI8XG, SX-FI24HF, SX-FI24GPP, and SX-F!48GPP modules running software release 07.3.00 or later.
- The port to which the cable is connected must be enabled when you issue the command to diagnose the cable. If the port is disabled, the command is rejected.
- If the port is operating at 100 Mbps half-duplex, the TDR test on one pair will fail.

- If the remote pair is set to forced 100 Mbps, any change in MDI/MDIX may cause the device to interpret the Multilevel Threshold-3 (MLT-3) as a reflected pulse, in which case, the device will report a faulty condition. In this scenario, it is recommended that you run the TDR test a few times for accurate results.

Virtual cable testing command syntax

To diagnose a cable using TDR, enter commands such as the following at the Privileged EXEC level of the CLI.

```
Brocade#phy cable-diag tdr 1
```

The above command diagnoses the cable attached to port 1.

When you issue the **phy-cable-diag** command, the command brings the port down for a second or two, then immediately brings the port back up.

Syntax: **phy cable-diag tdr** *port*

Viewing the results of the cable analysis

To display the results of the cable analysis, enter a command such as the following at the Privileged EXEC level of the CLI.

```
Brocade>show cable-diag tdr 1
Port      Speed Local pair Pair Length Remote pair Pair status
--------- ----- ---------- ----------- ----------- -----------
01        1000M Pair A     <50M        Pair B      Terminated
                Pair B     <50M        Pair A      Terminated
                Pair C     <50M        Pair D      Terminated
                Pair D     <50M        Pair C      Terminated
```

In the above output, **Local pair** indicates the assignment of wire pairs from left to right, where Pair A is the left-most pair. Table 18 shows the **Local pair** mapping to the T568A pin/pair and color assignment from the TIA/EIA-568-B standard.

TABLE 18 Local pair definition

Local pair	T568A pair and color assignment
Pair A	Pair 3 (green)
Pair B	Pair 2 (orange)
Pair C	Pair 1 (blue)
Pair D	Pair 4 (brown)

Figure 7 illustrates the T568A pin/pair assignment.

FIGURE 7 T568A pin/pair assignment

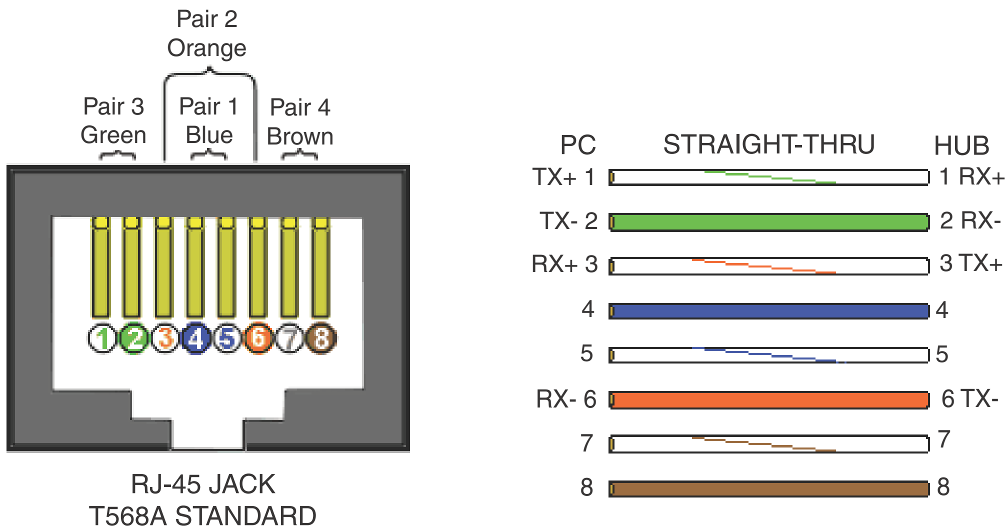

Syntax: **show cable-diag tdr** *port*

Table 19 defines the fields shown in the command output.

TABLE 19 Cable statistics

This line...	Displays...
Port	The port that was tested.
Speed	The port current line speed.
Local pair	The local link name. Refer to Table 18.
Pair Length	The cable length when terminated, or the distance to the point of fault when the line is not up.
Remote pair	The remote link name.
Pair status	The status of the link. This field displays one of the following: • Terminated: The link is up. • Shorted: A short is detected in the cable. • Open: An opening is detected in the cable. • ImpedMis: The impedance is mismatched. • Failed: The TDR test failed.

Digital optical monitoring

You can configure your Brocade device to monitor optical transceivers in the system, either globally or by specified ports. When this feature is enabled, the system will monitor the temperature and signal power levels for the optical transceivers in the specified ports. Console messages and Syslog messages are sent when optical operating conditions fall below or rise above the XFP, SFP, and SFP+ manufacturer recommended thresholds.

Digital optical monitoring configuration limitations

A Brocade chassis device can monitor a maximum of 24 SFPs and 12 XFPs.

Enabling digital optical monitoring

To enable optical monitoring on all Brocade-qualified optics installed in the device, use the following command.

```
Brocade(config)#optical-monitor
```

To enable optical monitoring on a specific port, use the following command.

```
Brocade(config)#interface ethernet 1/1
Brocade(config-if-e10000-1/1)#optical-monitor
```

To enable optical monitoring on a range of ports, use the following command.

```
Brocade(config)#interface ethernet 1/1 to 1/2
Brocade(config-mif-e10000-1/1-1/2)#optical-monitor
```

Syntax: [**no**] **optical-monitor**

Use the **no** form of the command to disable digital optical monitoring.

Setting the alarm interval

You can optionally change the interval between which alarms and warning messages are sent. The default interval is three minutes. To change the interval, use the following command.

```
Brocade(config)#interface ethernet 1/1 to 1/2
Brocade(config-mif-e10000-1/1-1/2)#optical-monitor 10
```

Syntax: [**no**] **optical-monitor** [*alarm-interval*]

For *alarm-interval*, enter a value between 1 and 65535. Enter 0 to disable alarms and warning messages.

NOTE
The commands **no optical-monitor** and **optical-monitor 0** perform the same function. That is, they both disable digital optical monitoring.

Displaying information about installed media

Use the **show media**, **show media slot**, and **show media ethernet** commands to obtain information about the media devices installed per device, per slot, and per port. The results displayed from these commands provide the Type, Vendor, Part number, Version and Serial number of the SFP, SFP+, or XFP optical device installed in the port. If there is no SFP, SFP+, or XFP optical device installed in a port, the “Type” field will display “EMPTY”.

On ICX 6430 and ICX 6450 devices, 1G copper ports will always be shown with the type as 1G M-C (Gig-Copper), even if the ports are not connected.

Use the **show media** command to obtain information about the media devices installed in a device.

```
Brocade#show media
Port 1/1/1:  Type : 1G M-C (Gig-Copper)
Port 1/1/2:  Type : 1G M-C (Gig-Copper)
Port 1/1/3:  Type : 1G M-C (Gig-Copper)
Port 1/1/4:  Type : 1G M-C (Gig-Copper)
Port 1/1/5:  Type : 1G M-C (Gig-Copper)
Port 1/1/6:  Type : 1G M-C (Gig-Copper)
Port 1/1/7:  Type : 1G M-C (Gig-Copper)
Port 1/1/8:  Type : 1G M-C (Gig-Copper)
Port 1/1/9:  Type : 1G M-C (Gig-Copper)
Port 1/1/10:  Type : 1G M-C (Gig-Copper)
Port 1/1/11:  Type : 1G M-C (Gig-Copper)
Port 1/1/12:  Type : 1G M-C (Gig-Copper)
Port 1/1/13:  Type : 1G M-C (Gig-Copper)
Port 1/1/14:  Type : 1G M-C (Gig-Copper)
Port 1/1/15:  Type : 1G M-C (Gig-Copper)
Port 1/1/16:  Type : 1G M-C (Gig-Copper)
Port 1/1/17:  Type : 1G M-C (Gig-Copper)
Port 1/1/18:  Type : 1G M-C (Gig-Copper)
Port 1/1/19:  Type : 1G M-C (Gig-Copper)
Port 1/1/20:  Type : 1G M-C (Gig-Copper)
Port 1/1/21:  Type : 1G M-C (Gig-Copper)
Port 1/1/22:  Type : 1G M-C (Gig-Copper)
Port 1/1/23:  Type : 1G M-C (Gig-Copper)
Port 1/1/24:  Type : 1G M-C (Gig-Copper)
Port 1/2/1:  Type : 10GE SR 300m (SFP +)
Port 1/2/2:  Type : EMPTY
Port 1/2/3:  Type : 1G Twinax   1m (SFP)
Port 1/2/4:  Type : 1G Twinax   1m (SFP)
```

Use the **show media slot** command to obtain information about the media device installed in a slot.

```
Brocade#show media slot 1
Port    1/1: Type  : 1G M-SX(SFP)
             Vendor: Brocade Communications, Inc.   Version:
             Part# : PL-XPL-VC-S13-19   Serial#: 425HC109
Port    1/2: Type  : 1G M-SX(SFP)
             Vendor: Brocade Communications, Inc.   Version:
             Part# : PL-XPL-VC-S13-19   Serial#: 411HC0AH
Port    1/3: Type  : EMPTY
Port    1/4: Type  : 1G M-SX(SFP)
             Vendor: Brocade Communications, Inc.   Version: X1
             Part# : FTRJ-8519-3        Serial#: H11654K
Port    1/5: Type  : EMPTY
Port    1/6: Type  : EMPTY
Port    1/7: Type  : 100M M-FX-IR(SFP)
```

```
               Vendor: Brocade Communications, Inc.   Version: A
               Part# : FTLF1323P1BTR-FD   Serial#: UCT000T
Port    1/8: Type  : EMPTY
Port    1/9: Type  : 100M M-FX-LR(SFP)
               Vendor: Brocade Communications, Inc.   Version: A
               Part# : FTLF1323P1BTL-FD   Serial#: UD3085J
Port   1/10: Type  : EMPTY
Port   1/11: Type  : 100M M-FX-SR(SFP)
               Vendor: Brocade Communications, Inc.   Version: A
               Part# : FTLF1217P2BTL-F1   Serial#: UCQ003J
Port   1/12: Type  : EMPTY
Port   1/13: Type  : 100M M-FX-IR(SFP)
               Vendor: Brocade Communications, Inc.   Version: A
               Part# : FTLF1323P1BTR-F1   Serial#: PCA2XC5
```

Use the **show media ethernet** command to obtain information about the media device installed in a port.

```
Brocade#show media e 1/17
Port   1/17: Type  : 1G M-SX(SFP)
             Vendor: Brocade Communications, Inc.   Version:
             Part# : PL-XPL-VC-S13-19   Serial#: 425HC109
```

Syntax: **show media** [**slot** *slot-num* | **ethernet** [*slot-num/*]*port-num*]

Viewing optical monitoring information

You can view temperature and power information for qualified XFPs, SFPs, and SFP+ installed in a FastIron device.

Use the **show optic** *port-number* command to view information about an XFP, SFP, or SFP+ installed in a particular port. The following shows example output.

Optical monitoring feature will not work in the following scenarios:

- The port is DOWN.
- The port is configured as a stacking port.
- The the optic module does not support optical monitoring.
- For ICX 6430 devices only:
  - If an SFP+ optic is inserted in an SFP only port, the optic will not initialize.
  - If an SFP optic is inserted in an SFP+ only port, the optic will not initialize.
  - If an optic is inserted into a device that supports both SFP and SFP+ optics, use the **speed-duplex** command to set the port speed correctly.

```
Brocade#show optic 13
Port  Temperature   Tx Power     Rx Power     Tx Bias Current
+----+-----------+----------+------------+-------------------+
13   33.2968 C  -005.4075 dBm -007.4328 dBm    6.306 mA
       Normal       Normal        Normal      Normal
```

Syntax: **show optic** *port-number*

Use the **show optic slot** *slot-number* on a FastIron X Series chassis to view information about all qualified XFPs, SFPs, and SFP+ in a particular slot. The following shows example output.

```
Brocade>show optic slot 4
Port  Temperature   Tx Power     Rx Power     Tx Bias Current
+----+-----------+----------+------------+-------------------+
4/1   30.8242 C  -001.8822 dBm -002.5908 dBm   41.790 mA
        Normal       Normal         Normal          Normal
4/2   31.7070 C  -001.4116 dBm -006.4092 dBm   41.976 mA
        Normal       Normal         Normal          Normal
4/3   30.1835 C                -000.5794 dBm    0.000 mA
        Normal      Low-Alarm      Normal        Low-Alarm
4/4    0.0000 C                                 0.000 mA
        Normal       Normal         Normal          Normal
```

Syntax: **show optic slot** *slot-number*

NOTE
The **show optic slot** *slot-number* command is supported on the FSX 800 and FSX 1600 only.

NOTE
The **show optic** function takes advantage of information stored and supplied by the manufacturer of the XFP, SFP, or SFP+ transceiver. This information is an optional feature of the Multi-Source Agreement standard defining the optical interface. Not all component suppliers have implemented this feature set. In such cases where the XFP, SFP, or SFP+ transceiver does not supply the information, a "Not Available" message will be displayed for the specific port on which the module is installed.

The following table describes the information displayed by the **show optic** command.

TABLE 20 Output from the show optic command

Field	Description
Port	The Brocade port number.
Temperature	• The operating temperature, in degrees Celsius, of the optical transceiver. • The alarm status, as described in Table 21.
Tx Power	• The transmit power signal, in decibels (dB), of the measured power referenced to one milliwatt (mW). • The alarm status, as described in Table 21.
Rx Power	• The receive power signal, in decibels (dB), of the measured power referenced to one milliwatt (mW). • The alarm status, as described in Table 21
Tx Bias Current	• The transmit bias power signal, in milliamperes (mA). • The alarm status, as described in Table 21.

For Temperature, Tx Power, Rx Power, and Tx Bias Current in the **show optic** command output, values are displayed along with one of the following alarm status values: Low-Alarm, Low-Warn, Normal, High-Warn or High-Alarm. The thresholds that determine these status values are set by the manufacturer of the optical transceivers. Table 21 describes each of these status values.

TABLE 21 Alarm status value description

Status value	Description
Low-Alarm	Monitored level has dropped below the "low-alarm" threshold set by the manufacturer of the optical transceiver.
Low-Warn	Monitored level has dropped below the "low-warn" threshold set by the manufacturer of the optical transceiver.
Normal	Monitored level is within the "normal" range set by the manufacturer of the optical transceiver.
High-Warn	Monitored level has climbed above the "high-warn" threshold set by the manufacturer of the optical transceiver.
High-Alarm	Monitored level has climbed above the "high-alarm" threshold set by the manufacturer of the optical transceiver.

Viewing optical transceiver thresholds

The thresholds that determine the alarm status values for an optical transceiver are set by the manufacturer of the XFP, SFP, or SFP+. To view the thresholds for a qualified optical transceiver in a particular port, use the **show optic threshold** command as shown below.

```
Brocade>show optic threshold 2/2
Port 2/2 sfp monitor thresholds:
Temperature High alarm                5a00          90.0000 C
Temperature Low alarm                 d300         -45.0000 C
Temperature High warning              5500          85.0000 C
Temperature Low warning               d800         -40.0000 C
Supply Voltage High alarm             9088
Supply Voltage Low alarm              7148
Supply Voltage High warning           8ca0
Supply Voltage Low warning            7530
TX Bias High alarm                    7530          60.000 mA
TX Bias Low alarm                     01f4           1.000 mA
TX Bias High warning                  61a8          50.000 mA
TX Bias Low warning                   05dc           3.000 mA
TX Power High alarm                   1f07        -001.0001 dBm
TX Power Low alarm                    02c4        -011.4996 dBm
TX Power High warning                 18a6        -001.9997 dBm
TX Power Low warning                  037b        -010.5012 dBm
RX Power High alarm                   2710         000.0000 dBm
RX Power Low alarm                    0028        -023.9794 dBm
RX Power High warning                 1f07        -001.0001 dBm
RX Power Low warning                     0032         -023.0102 dBm
```

Syntax: **show optic threshold** *port*

For Temperature, Supply Voltage, TX Bias, TX Power, and RX Power, values are displayed for each of the following four alarm and warning settings: High alarm, Low alarm, High warning, and Low warning. The hexadecimal values are the manufacturer internal calibrations, as defined in the SFF-8472 standard. The other values indicate at what level (above the high setting or below the low setting) the system should send a warning message or an alarm. Note that these values are set by the manufacturer of the optical transceiver, and cannot be configured.

Syslog messages for optical transceivers

The system generates Syslog messages for optical transceivers in the following circumstances:

- The temperature, supply voltage, TX Bias, TX power, or TX power value goes above or below the high or low warning or alarm threshold set by the manufacturer.
- The optical transceiver does not support digital optical monitoring.
- The optical transceiver is not qualified, and therefore not supported by Brocade.

For details about the above Syslog messages, refer to Appendix A, “Syslog messages”.

Chapter

Syslog

10

Table 22 lists individual Brocade switches and the Syslog features they support. These features are supported in the Layer 2 and Layer 3 software images, except where explicitly noted.

TABLE 22 Supported Syslog features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6430	ICX 6450
Syslog messages	Yes	Yes	Yes	Yes	Yes
Real-time display of Syslog messages	Yes	Yes	Yes	Yes	Yes
Real-time display for Telnet or SSH sessions	Yes	Yes	Yes	Yes	Yes
Show log on all terminals	Yes	Yes	Yes	Yes	Yes
Time stamps	Yes	Yes	Yes	Yes	Yes
Multiple Syslog server logging (up to 6 Syslog servers)	Yes	Yes	Yes	Yes	Yes
Disabling logging of a message level	Yes	Yes	Yes	Yes	Yes
Changing the number of entries the local buffer can hold	Yes	Yes	Yes	Yes	Yes
Changing the log facility	Yes	Yes	Yes	Yes	Yes
Displaying Interface names in Syslog messages	Yes	Yes	Yes	Yes	Yes
Displaying TCP and UDP port numbers in Syslog messages	Yes	Yes	Yes	Yes	Yes
Retaining Syslog messages after a soft reboot	Yes	Yes	Yes	Yes	Yes
Clearing Syslog messages from the local buffer	Yes	Yes	Yes	Yes	Yes
Syslog messages for hardware errors	Yes	No	No	No	No
Default Log Buffer Size	Yes	Yes	Yes	Yes	Yes

This chapter contains the following sections:

This chapter describes how to display Syslog messages and how to configure the Syslog facility, and lists the Syslog messages that Brocade devices can display during standard operation. Refer to “Syslog” on page 243 for a list of Syslog messages.

About Syslog messages

Brocade software can write syslog messages to provide information at the following severity levels:

- Emergencies
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational
- Debugging

The device writes the messages to a local buffer.

You also can specify the IP address or host name of up to six Syslog servers. When you specify a Syslog server, the Brocade device writes the messages both to the system log and to the Syslog server.

Using a Syslog server ensures that the messages remain available even after a system reload. The Brocade local Syslog buffer is cleared during a system reload or reboot, but the Syslog messages sent to the Syslog server remain on the server.

NOTE
To enable the Brocade device to retain Syslog messages after a soft reboot (**reload** command). Refer to "Retaining Syslog messages after a soft reboot" on page 254.

The Syslog service on a Syslog server receives logging messages from applications on the local host or from devices such as a Layer 2 Switch or Layer 3 Switch. Syslog adds a time stamp to each received message and directs messages to a log file. Most Unix workstations come with Syslog configured. Some third party vendor products also provide Syslog running on NT.

Syslog uses UDP port 514 and each Syslog message thus is sent with destination port 514. Each Syslog message is one line with Syslog message format. The message is embedded in the text portion of the Syslog format. There are several subfields in the format. Keywords are used to identify each subfield, and commas are delimiters. The subfield order is insensitive except that the text subfield should be the last field in the message. All the subfields are optional.

Displaying Syslog messages

To display the Syslog messages in the device local buffer, enter the **show logging** command at any level of the CLI. The following shows an example display output.

```
Brocade>#show logging

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
   Buffer logging: level ACDMEINW, 3 messages logged
   level code: A=alert C=critical D=debugging M=emergency E=error
            I=informational N=notification W=warning

Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
```

```
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

For information about the Syslog configuration information, time stamps, and dynamic and static buffers, refer to "Displaying the Syslog configuration" on page 246.

Enabling real-time display of Syslog messages

By default, to view Syslog messages generated by a Brocade device, you need to display the Syslog buffer or the log on a Syslog server used by the Brocade device.

You can enable real-time display of Syslog messages on the management console. When you enable this feature, the software displays a Syslog message on the management console when the message is generated. However, to enable display of real-time Syslog messages in Telnet or SSH sessions, you also must enable display within the individual sessions.

To enable real-time display of Syslog messages, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)#logging console
```

Syntax: [**no**] **logging console**

This command enables the real-time display of Syslog messages on the serial console. You can enter this command from the serial console or a Telnet or SSH session.

Enabling real-time display for a Telnet or SSH session

To also enable the real-time display for a Telnet or SSH session, enter the following command from the Privileged EXEC level of the session.

```
telnet@Brocade#terminal monitor
Syslog trace was turned ON
```

Syntax: **terminal monitor**

Notice that the CLI displays a message to indicate the status change for the feature. To disable the feature in the management session, enter the **terminal monitor** command again. The command toggles the feature on and off.

```
telnet@Brocade#terminal monitor
Syslog trace was turned OFF
```

Here is an example of how the Syslog messages are displayed.

```
telnet@Brocade#terminal monitor
Syslog trace was turned ON
SYSLOG: <9>Brocade, Power supply 2, power supply on left connector, failed

SYSLOG: <14>Brocade, Interface ethernet 6, state down

SYSLOG: <14>Brocade, Interface ethernet 2, state up
```

Displaying real-time Syslog messages

Any terminal logged on to a Brocade switch can receive real-time Syslog messages when the **terminal monitor** command is issued.

Syslog service configuration

The procedures in this section describe how to perform the following Syslog configuration tasks:

- Specify a Syslog server. You can configure the Brocade device to use up to six Syslog servers. (Use of a Syslog server is optional. The system can hold up to 1000 Syslog messages in an internal buffer.)
- Change the level of messages the system logs.
- Change the number of messages the local Syslog buffer can hold.
- Display the Syslog configuration.
- Clear the local Syslog buffer.

Logging is enabled by default, with the following settings:

- Messages of all severity levels (Emergencies – Debugging) are logged.
- By default, up to 50 messages are retained in the local Syslog buffer. This can be changed.
- No Syslog server is specified.

Displaying the Syslog configuration

To display the Syslog parameters currently in effect on a Brocade device, enter the following command from any level of the CLI.

```
Brocade>#show logging

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 3 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning

Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed

Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 1/4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

Syntax: **show logging**

The Syslog display shows the following configuration information, in the rows above the log entries themselves.

TABLE 23 CLI display of Syslog buffer configuration

Field	Definition
Syslog logging	The state (enabled or disabled) of the Syslog buffer.
messages dropped	The number of Syslog messages dropped due to user-configured filters. By default, the software logs messages for all Syslog levels. You can disable individual Syslog levels, in which case the software filters out messages at those levels. Refer to "Disabling logging of a message level" on page 251. Each time the software filters out a Syslog message, this counter is incremented.
flushes	The number of times the Syslog buffer has been cleared by the **clear logging** command. Refer to "Clearing the Syslog messages from the local buffer" on page 254.
overruns	The number of times the dynamic log buffer has filled up and been cleared to hold new entries. For example, if the buffer is set for 100 entries, the 101st entry causes an overrun. After that, the 201st entry causes a second overrun.
level	The message levels that are enabled. Each letter represents a message type and is identified by the key (level code) below the value. If you disable logging of a message level, the code for that level is not listed.
messages logged	The total number of messages that have been logged since the software was loaded.
level code	The message levels represented by the one-letter codes.

Static and dynamic buffers

The software provides two buffers:

- Static – logs power supply failures, fan failures, and temperature warning or shutdown messages
- Dynamic – logs all other message types

In the static log, new messages replace older ones, so only the most recent message is displayed. For example, only the most recent temperature warning message will be present in the log. If multiple temperature warning messages are sent to the log, the latest one replaces the previous one. The static buffer is not configurable.

The message types that appear in the static buffer do not appear in the dynamic buffer. The dynamic buffer contains up to the maximum number of messages configured for the buffer (50 by default), then begins removing the oldest messages (at the bottom of the log) to make room for new ones.

The static and dynamic buffers are both displayed when you display the log.

```
Brocade#show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 3 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
          I=informational N=notification W=warning

Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dec 15 19:00:14:A:Fan 2, fan on left connector, failed

Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

Notice that the static buffer contains two separate messages for fan failures. Each message of each type has its own buffer. Thus, if you replace fan 1 but for some reason that fan also fails, the software replaces the first message about the failure of fan 1 with the newer message. The software does not overwrite the message for fan 2, unless the software sends a newer message for fan 2.

Clearing log entries

When you clear log entries, you can selectively clear the static or dynamic buffer, or you can clear both. For example, to clear only the dynamic buffer, enter the following command at the Privileged EXEC level.

```
Brocade#clear logging dynamic-buffer
```

Syntax: **clear logging** [**dynamic-buffer** | **static-buffer**]

You can specify **dynamic-buffer** to clear the dynamic buffer or **static-buffer** to clear the static buffer. If you do not specify a buffer, both buffers are cleared.

Time stamps

The contents of the time stamp differ depending on whether you have set the time and date on the onboard system clock:

- If you have set the time and date on the onboard system clock, the date and time are shown in the following format.

  mm dd hh:mm:ss

  where

  - *mm* – abbreviation for the name of the month
  - *dd* – day
  - *hh* – hours
  - *mm* – minutes
  - *ss* – seconds

  For example, "Oct 15 17:38:03" means October 15 at 5:38 PM and 3 seconds.

- If you have not set the time and date on the onboard system clock, the time stamp shows the amount of time that has passed since the device was booted, in the following format.

  num d *num* h *num* m *num* s

  where

  - *num* d – day
  - *num* h – hours
  - *num* m – minutes
  - *num* s – seconds

  For example, “188d1h01m00s” means the device had been running for 188 days, 11 hours, one minute, and zero seconds when the Syslog entry with this time stamp was generated.

Example of Syslog messages on a device with the onboard clock set

The example shows the format of messages on a device where the onboard system clock has been set. Each time stamp shows the month, the day, and the time of the system clock when the message was generated. For example, the system time when the most recent message (the one at the top) was generated was October 15 at 5:38 PM and 3 seconds.

```
Brocade#show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 38 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning

Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dec 15 19:00:14:A:Fan 2, fan on left connector, failed

Dynamic Log Buffer (50 entries):
Oct 15 17:38:03:warning:list 101 denied tcp 10.157.22.191(0)(Ethernet 18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
Oct 15 07:03:30:warning:list 101 denied tcp 10.157.22.26(0)(Ethernet 18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
Oct 15 06:58:30:warning:list 101 denied tcp 10.157.22.198(0)(Ethernet 18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
```

Example of Syslog messages on a device wih the onboard clock not set

The example shows the format of messages on a device where the onboard system clock is not set. Each time stamp shows the amount of time the device had been running when the message was generated. For example, the most recent message, at the top of the list of messages, was generated when the device had been running for 21 days, seven hours, two minutes, and 40 seconds.

```
Brocade#show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 38 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning

Static Log Buffer:

Dynamic Log Buffer (50 entries):
21d07h02m40s:warning:list 101 denied tcp 10.157.22.191(0)(Ethernet 4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
19d07h03m30s:warning:list 101 denied tcp 10.157.22.26(0)(Ethernet 4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
17d06h58m30s:warning:list 101 denied tcp 10.157.22.198(0)(Ethernet 4/18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
```

Disabling or re-enabling Syslog

Syslog is enabled by default. To disable it, enter the **logging on** command at the global CONFIG level.

```
Brocade(config)#no logging on
```

Syntax: [**no**] **logging on** [*udp-port*]

The *udp-port* parameter specifies the application port used for the Syslog facility. The default is 514.

To re-enable logging, re-enter the **logging on** command.

```
Brocade(config)#logging on
```

This command enables local Syslog logging with the following defaults:

- Messages of all severity levels (Emergencies – Debugging) are logged.
- Up to 50 messages are retained in the local Syslog buffer.
- No Syslog server is specified.

Specifying a Syslog server

To specify a Syslog server, enter the **logging host** command.

```
Brocade(config)#logging host 10.0.0.99
```

Syntax: **logging host** *ip-addr* | *server-name*

Specifying an additional Syslog server

To specify an additional Syslog server, enter the **logging host** *ip-addr* command again. You can specify up to six Syslog servers.

```
Brocade(config)#logging host 10.0.0.99
```

Syntax: **logging host** *ip-addr* | *server-name*

Disabling logging of a message level

To change the message level, disable logging of specific message levels. You must disable the message levels on an individual basis.

For example, to disable logging of debugging and informational messages, enter the following commands.

```
Brocade(config)#no logging buffered debugging
Brocade(config)#no logging buffered informational
```

Syntax: [**no**] **logging buffered** *level* | *num-entries*

The *level* parameter can have one of the following values:

- alerts
- critical
- debugging
- emergencies
- errors
- informational
- notifications
- warnings

The commands in the example above change the log level to notification messages or higher. The software will not log informational or debugging messages. The changed message level also applies to the Syslog servers.

Changing the number of entries the local buffer can hold

You also can use the **logging buffered** command to change the number of entries the local Syslog buffer can store. For example.

```
Brocade(config)#logging buffered 1000
Brocade(config)#write memory
Brocade(config)#exit
Brocade#reload
```

Syntax: [no] **logging buffered** *num*

The default number of messages is 50. For FastIron devices, you can set the Syslog buffer limit from 1 – 1000 entries.

Local buffer configuration notes

- You must save the configuration and reload the software to place the change into effect.
- If you decrease the size of the buffer, the software clears the buffer before placing the change into effect.
- If you increase the size of the Syslog buffer, the software will clear some of the older locally buffered Syslog messages.

Changing the log facility

The Syslog daemon on the Syslog server uses a facility to determine where to log the messages from the Brocade device. The default facility for messages the Brocade device sends to the Syslog server is “user”. You can change the facility using the following command.

NOTE
You can specify only one facility. If you configure the Brocade device to use two Syslog servers, the device uses the same facility on both servers.

```
Brocade(config)#logging facility local0
```

Syntax: **logging facility** *facility-name*

The *facility-name* can be one of the following:

- kern – kernel messages
- user – random user-level messages
- mail – mail system
- daemon – system daemons
- auth – security or authorization messages
- syslog – messages generated internally by Syslog
- lpr – line printer subsystem
- news – netnews subsystem
- uucp – uucp subsystem
- sys9 – cron/at subsystem
- sys10 – reserved for system use
- sys11 – reserved for system use
- sys12 – reserved for system use
- sys13 – reserved for system use
- sys14 – reserved for system use
- cron – cron/at subsystem
- local0 – reserved for local use
- local1 – reserved for local use
- local2 – reserved for local use
- local3 – reserved for local use
- local4 – reserved for local use
- local5 – reserved for local use
- local6 – reserved for local use
- local7 – reserved for local use

Displaying interface names in Syslog messages

By default, an interface slot number (if applicable) and port number are displayed when you display Syslog messages. If you want to display the name of the interface instead of its number, enter the following command:

```
FastIron(config)# ip show-portname
```

This command is applied globally to all interfaces on Layer 2 Switches and Layer 3 Switches.

Syntax: [no] Ip show-portname

By default, Syslog messages show the interface type, such as “ethernet”, and so on. For example, you see the following

```
SYSLOG: <14>0d00h02m18s:ICX6610-48P Router System: Interface ethernet 1/1/5,
state up
```

However, if ip show-portname is configured and a name has been assigned to the port, the port name replaces the interface type as in the example below, where “port5_name” is the name of the port.

```
SYSLOG: <14>0d00h02m18s:ICX6610-48P Router System: Interface port5_name 1/1/5,
state up
```

Also, when you display the messages in the Syslog, you see the interface name under the Dynamic Log Buffer section. The actual interface number is appended to the interface name. For example, if the interface name is "lab" and its port number is "2", you see "lab2" displayed as in the example below:

```
Brocade# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Buffer logging: level ACDMEINW, 3 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
          I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet Lab2, state up
Dec 15 18:45:15:I:Warm start
```

Displaying TCP or UDP port numbers in Syslog messages

The command ip show-service-number-in-log allows you to change the display of TCP or UDP application information from the TCP or UDP well-known port name to the TCP or UDP port number. For example, when this command is in effect, the Brocade device will display http (the well-known port name) instead of 80 (the port number) in the output of show commands, and other commands that contain application port information. By default, Brocade devices display TCP or UDP application information in named notation.

To display TCP or UDP port numbers instead of their names, enter the following command.

```
Brocade(config)#ip show-service-number-in-log
```

Syntax: [no] ip show-service-number-in-log

Retaining Syslog messages after a soft reboot

You can configure the device to save the System log (Syslog) after a soft reboot (**reload** command).

Syslog reboot configuration considerations

- If the Syslog buffer size was set to a different value using the CLI command **logging buffered**, the System log will be cleared after a soft reboot, even when this feature (**logging persistence**) is in effect. This will occur only with a soft reboot immediately following a Syslog buffer size change. A soft reboot by itself will not clear the System log. To prevent the system from clearing the System log, leave the number of entries allowed in the Syslog buffer unchanged.
- This feature does not save Syslog messages after a hard reboot. When the Brocade device is power-cycled, the Syslog messages are cleared.
- If **logging persistence** is enabled and you load a new software image on the device, you must first clear the log if you want to reload the device. (Refer to "Clearing the Syslog messages from the local buffer" on page 254.)

To configure the device to save the System log messages after a soft reboot, enter the following command.

```
Brocade(config)#logging persistence
```

Syntax: [**no**] **logging persistence**

Enter **no logging persistence** to disable this feature after it has been enabled.

Clearing the Syslog messages from the local buffer

To clear the Syslog messages stored in the local buffer of the Brocade device, enter the **clear logging** command.

```
Brocade#clear logging
```

Syntax: **clear logging**

Syslog messages for hardware errors

NOTE
This feature is supported on FastIron X Series devices only. It is **not** supported on FCX and ICX devices.

FastIron Chassis devices support the display of hardware read and write errors encountered on a slot or module during bootup and during normal system operations. There are four types of errors, which may cause the system to disable or power down the modules on which they occur:

- Configuration read error
- Configuration write error
- Memory read error
- Memory write error

The following shows examples of some hardware errors in the **show logging** display output.

```
Brocade>#show logging

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
   Buffer logging: level ACDMEINW, 3 messages logged
   level code: A=alert C=critical D=debugging M=emergency E=error
          I=informational N=notification W=warning

Dynamic Log Buffer (50 lines):
0d00h00m27s:I:System: Interface ethernet mgmt1, state up
0d00h00m26s:N:powered On switch Fabric
0d00h00m17s:N:powered On switch Fabric
0d00h00m08s:I:System: Warm start
0d00h00m08s:I:SNMP: read-only community added by  from  session
0d00h00m02s:A:System: Module in slot 5 encountered unrecoverable PCI bridge
validation failure. Module will be deleted.
0d00h00m02s:A:System: Module in slot 5 encountered unrecoverable PCI config read
failure. Module will be deleted.
0d00h00m02s:A:System: Module in slot 5 encountered PCI config read error: Bus 10,
Dev 3, Reg Offset 0.
0d00h00m00s:W:System: Fan speed changed automatically to 1
```

Syslog messages (alerts) for hardware errors are listed in Table 1 on page 349.

Chapter 11

Network Monitoring

Table 24 lists the individual FastIron switches and the network monitoring features they support. These features are supported in the Layer 2 and Layer 3 software images, except where explicitly noted.

TABLE 24 Supported network monitoring features

Feature	FSX 800 FSX 1600	FCX	ICX 6610	ICX 6430	ICX 6450
Traffic counters for outbound traffic	Yes	No	No	Yes	Yes
Egress queue counters	No	Yes	Yes	No	No
Remote monitoring (RMON)	Yes	Yes	Yes	Yes	Yes
Specifying the maximum number of entries allowed in the RMON Control Table	Yes	Yes	Yes	Yes	Yes
sFlow version 2	Yes	Yes	Yes	No	Yes
sFlow version 5 (default)	Yes	Yes	Yes	No	Yes
sFlow support for IPv6 packets	Yes	Yes	Yes	No	Yes
Uplink utilization lists	Yes	Yes	Yes	Yes	Yes

This chapter contains the following sections:

Basic system management

The following sections contain procedures for basic system management tasks.

Viewing system information

You can access software and hardware specifics for a Brocade Layer 2 Switch or Layer 3 Switch. For software specifics, refer to "Software versions installed and running on a device" on page 72.

To view the software and hardware details for the system, enter the **show version** command. The following shows an example output.

```
Brocade#show version

==========================================================================
Active Management CPU [Slot-9]:
  SW: Version 04.3.00b17T3e3 Copyright (c) 1996-2008 Brocade Communications,
Inc., Inc.
      Compiled on Sep 25 2008 at 04:09:20 labeled as SXR04300b17
      (4031365 bytes) from Secondary sxr04300b17.bin
      BootROM: Version 04.0.00T3e5 (FEv2)
  HW: ANR-Chassis FastIron SX 1600-PREM (PROM-TYPE SX-FIL3U)
      Serial  #: TExxxxxxxx
==========================================================================
SL 3: SX-FI424C 24-port Gig Copper
      Serial  #: CYxxxxxxxxx
      P-ASIC  4: type 00D1, rev D2  subrev 00
      P-ASIC  5: type 00D1, rev D2  subrev 00
==========================================================================
SL 9: SX-FI8GMR4 8-port Management
      Serial  #: CHxxxxxxxx
      P-ASIC 16: type 00D1, rev D2  subrev 00
==========================================================================
SL 14: SX-FI42XGW 2-port 10G LAN/WAN
      Serial  #: Invalid
      P-ASIC 26: type 01D1, rev 00  subrev 00
      P-ASIC 27: type 01D1, rev 00  subrev 00
==========================================================================
Active Management Module:
  660 MHz Power PC processor 8541 (version 32/0020) 66 MHz bus
  512 KB boot flash memory
16384 KB code flash memory
  512 MB DRAM
The system uptime is 2 minutes 13 seconds
The system : started=warm start   reloaded=by "reload"
*** NOT FOR PRODUCTION ***
*** AUTO SHUTDOWN IS OFF. PLEASE ACTIVATE WITH auto-shutdown ***
```

The following hardware details are listed in the output of the **show version** command:

- Chassis type
- PROM type (if applicable)
- Chassis serial number
- Management and interface module serial numbers and ASIC types

For a description of the software details in the output of the **show version** command, refer to "Software versions installed and running on a device" on page 72.

Syntax: show version

Viewing configuration information

You can view a variety of configuration details and statistics with the **show** option. The **show** option provides a convenient way to check configuration changes before saving them to flash.

The show options available will vary for Layer 2 Switches and Layer 3 Switches and by configuration level.

To determine the available show commands for the system or a specific level of the CLI, enter the following command.

```
Brocade#show ?
```

Syntax: show *option*

You also can enter “show” at the command prompt, then press the TAB key.

Viewing port statistics

Port statistics are polled by default every 10 seconds.

You can view statistics for ports by entering the following **show** commands:

- show interfaces
- show configuration
- show statistics

To display the statistics, enter a command such as the following.

```
Brocade#show statistics ethernet 1/3
Port  Link State     Dupl Speed Trunk Tag Priori MAC            Name
1/3     Up   Forward   Half 100M  None  No  level0 0000.0000.0102

 Port 1/3 Counters:
         InOctets                 3200          OutOctets                   256
           InPkts                   50            OutPkts                     4
  InBroadcastPkts                    0    OutBroadcastPkts                    3
  InMulticastPkts                   48    OutMulticastPkts                    0
    InUnicastPkts                    2      OutUnicastPkts                    1
        InBadPkts                    0
      InFragments                    0
       InDiscards                    0          OutErrors                     0
              CRC                    0         Collisions                     0
         InErrors                    0      LateCollisions                    0
      InGiantPkts                    0
      InShortPkts                    0
         InJabber                    0
   InFlowCtrlPkts                    0     OutFlowCtrlPkts                    0
     InBitsPerSec                  264       OutBitsPerSec                   16
     InPktsPerSec                    0       OutPktsPerSec                    0
     InUtilization                 0.00%     OutUtilization                0.00%
```

Syntax: show statistics [**ethernet** [*port*]

Table 25 lists the statistics displayed in the output of the **show statistics** command.

TABLE 25 Port statistics

Parameter	Description
Port configuration	
Port	The port number.
Link	The link state.
State	The STP state.

TABLE 25 Port statistics (Continued)

Parameter	Description
Dupl	The mode (full-duplex or half-duplex).
Speed	The port speed (10M, 100M, or 1000M).
Trunk	The trunk group number, if the port is a member of a trunk group.
Tag	Whether the port is a tagged member of a VLAN.
Priori	The QoS forwarding priority of the port (level0 – level7).
MAC	The MAC address of the port.
Name	The name of the port, if you assigned a name.
Statistics	
InOctets	The total number of good octets and bad octets received.
OutOctets	The total number of good octets and bad octets sent.
InPkts	The total number of packets received. The count includes rejected and local packets that are not sent to the switching core for transmission.
OutPkts	The total number of good packets sent. The count includes unicast, multicast, and broadcast packets.
InBroadcastPkts	The total number of good broadcast packets received.
OutBroadcastPkts	The total number of good broadcast packets sent.
InMulticastPkts	The total number of good multicast packets received.
OutMulticastPkts	The total number of good multicast packets sent.
InUnicastPkts	The total number of good unicast packets received.
OutUnicastPkts	The total number of good unicast packets sent.
InBadPkts	The total number of packets received for which one of the following is true: • The CRC was invalid. • The packet was oversized. • Jabbers: The packets were longer than 1518 octets and had a bad FCS. • Fragments: The packets were less than 64 octets long and had a bad FCS. • The packet was undersized (short).
InFragments	The total number of packets received for which both of the following was true: • The length was less than 64 bytes. • The CRC was invalid.
InDiscards	The total number of packets that were received and then dropped due to a lack of receive buffers.
OutErrors	The total number of packets with internal transmit errors such as TX underruns.
CRC	The total number of packets received for which all of the following was true: • The data length was between 64 bytes and the maximum allowable frame size. • No Collision or Late Collision was detected. • The CRC was invalid.
Collisions	The total number of packets received in which a Collision event was detected.
InErrors	The total number of packets received that had Alignment errors or phy errors.
LateCollisions	The total number of packets received in which a Collision event was detected, but for which a receive error (Rx Error) event was not detected.

TABLE 25 **Port statistics (Continued)**

Parameter	Description
InGiantPkts	The total number of packets for which all of the following was true: • The data length was longer than the maximum allowable frame size. • No Rx Error was detected. **NOTE:** Packets are counted for this statistic regardless of whether the CRC is valid or invalid.
InShortPkts	The total number of packets received for which all of the following was true: • The data length was less than 64 bytes. • No Rx Error was detected. • No Collision or Late Collision was detected. **NOTE:** Packets are counted for this statistic regardless of whether the CRC is valid or invalid.
InJabber	The total number of packets received for which all of the following was true: • The data length was longer than the maximum allowable frame size. • No Rx Error was detected. • The CRC was invalid.
InFlowCtrlPkts	The total number of flow control packets received.
OutFlowCtrlPkts	The total number of flow control packets transmitted.
InBitsPerSec	The number of bits received per second.
OutBitsPerSec	The number of bits sent per second.
InPktsPerSec	The number of packets received per second.
OutPktsPerSec	The number of packets sent per second.
InUtilization	The percentage of the port bandwidth used by received traffic.
OutUtilization	The percentage of the port bandwidth used by sent traffic.

Viewing STP statistics

You can view a summary of STP statistics for Layer 2 Switches and Layer 3 Switches. STP statistics are by default polled every 10 seconds.

To view spanning tree statistics, enter the **show span** command. To view STP statistics for a VLAN, enter the **span vlan** command.

Clearing statistics

You can clear statistics for many parameters using the **clear** command.

To determine the available **clear** commands for the system, enter the **clear** command at the Privileged EXEC level of the CLI.

```
Brocade#clear ?
```

Syntax: clear *option*

You also can enter "clear" at the command prompt, then press the TAB key.

Traffic counters for outbound traffic

You can configure traffic counters (also called transmit counters) that enable the Brocade device to count the following packet types on a port or port region:

- broadcast packets
- multicast packets
- unicast packets
- dropped packets due to congestion and egress filtering

Depending on the parameters specified with the traffic counter configuration, traffic counters record the number of outbound packets from any combination of the following sources:

- a specific port or all ports in a specific port region
- a specific VLAN or all VLANs
- a specific 802.1p priority queue or all priority queues

Traffic counters configuration notes

Consider the following rules when configuring traffic counters for outbound traffic.

- This feature is supported on FastIron X Series devices only.
- This feature is supported in the Layer 2 and Layer 3 codes.
- This feature applies to physical ports only, including 10 Gbps Ethernet ports and trunk ports. It does not apply to virtual interfaces.
- Once the enhanced traffic counters are read using the **show transmit-counter values** command, the counters are cleared (reset to zero).
- For each port region, you can enable a maximum of two traffic counters, regardless of whether traffic counters are enabled on individual ports or on all ports in the port region.
- Traffic counters increase for bridged filtered outbound traffic when any of the following conditions occur:
  - The port is disabled or the link is down.
  - The port or port region does not belong to the VLAN specified in the transmit counter configuration.
  - A Layer 2 protocol (e.g., spanning tree) has the port in a Blocked state.
  - The source port needs to be suppressed for multi-target packets.
  - The priority queue specified in the traffic counter is not allowed for some other reason.

- Unknown unicast and unregistered multicast packets are filtered.

Traffic counters configuration syntax

This section provides the syntax and configuration examples for enhanced traffic counters.

Example

To configure traffic counters for outbound traffic on a specific port, enter a command such as the following.

```
Brocade(config)#transmit-counter 4 port 18 only vlan 1 prio 7 enable
```

The above command creates and enables traffic counter 4 on port 18. The device will count the number of packets sent out on port 18 that are in VLAN 1 and have a priority queue of 7.

Example

To configure traffic counters for outbound traffic in a specific port region, enter a command such as the following.

```
Brocade(config)#transmit-counter 1 port 1 region vlan all prio all enable
```

The above command creates and enables traffic counter 1 on all ports that are in the same port region as port 1. The device will count the number of packets transmitted in this port region that belong to any VLAN and have any assigned priority queue.

Syntax: [**no**] **transmit-counter** *counter-ID* **port** [*slotnum/*]*port-num* **only** | **region vlan** *vlan-ID* | **all priority** *priority-queue* | **all enable**

Enter the **no** form of the command to remove the outbound traffic counter.

The *counter-ID* parameter identifies the traffic counter. You can configure up to 64 traffic counters. Enter a number from 1 – 64.

The *slotnum* parameter is required on chassis devices.

The *port-num* parameter is the port number to which enhanced traffic counters will apply. Enter the port number followed by **only** to apply the enhanced traffic counter to a specific port, or enter the port number followed by **region** to apply the enhanced traffic counter to all of the ports in the port region.

The *vlan-ID* parameter identifies the VLAN ID for which outbound traffic will be counted. Enter a number from 0 – 4095 or enter **all** to indicate all VLANs.

The *priority-queue* parameter identifies the 802.1p priority queue for which traffic will be counted. Enter a number from 0 – 7 or enter **all** to indicate all priority queues.

Displaying enhanced traffic counter profiles

To display the details of the traffic counters configured on your device, enter the **show transmit-counter profiles** command. The following shows an example output.

```
Brocade#show transmit-counter profiles
Tx Counter      Port(s)   Vlan Id  Priority  Device    Set
         1    1 -   12       All       All   Dev 0   Set0
         4          18         1         7   Dev 1   Set0
        10   13 -   24       100       All   Dev 1   Set1
```

Displaying enhanced traffic counter statistics

To display the traffic counters for outbound traffic, enter the **show transmit-counter profiles** command.

NOTE
Once the enhanced traffic counters are displayed, the counters are cleared (reset to zero).

The following shows an example output.

```
Brocade#show transmit-counter values 1

Transmit Queue Counter Values for Counter 1:
Transmitted Frames:
 Known Unicast              : 17204
 Multicast & Unknown Unicast : 2797
 Broadcast                  : 5
Dropped Frames:
 Bridge Egress Filtered     : 2
 Congestion Drops           : 0

Brocade#show transmit-counter values 4

Transmit Queue Counter Values for Counter 4:
Transmitted Frames:
 Known Unicast              : 124
 Multicast & Unknown Unicast : 2752
 Broadcast                  : 0
Dropped Frames:
 Bridge Egress Filtered     : 37
 Congestion Drops           : 0
```

Syntax: **show transmit-counter values** *number*

where *number* identifies a valid enhanced traffic counter and is a value from 1 – 64.

TABLE 26 Outbound traffic counter statistics

This line...	Displays...
Transmitted frames	
Known Unicast	The number of known unicast packets transmitted.
Multicast & Unknown Unicast	The number of multicast and unknown unicast packets transmitted.
Broadcast	The number of broadcast packets transmitted.
Dropped Frames	

TABLE 26 Outbound traffic counter statistics (Continued)

This line...	Displays...
Bridge Egress Filtered	The number of bridged outbound packets that were filtered and dropped. This number includes the number of packets that were dropped because of any one of the following conditions: • The port was disabled or the link was down. • The port or port region does not belong to the VLAN specified in the transmit counter configuration. • A Layer 2 protocol (e.g., spanning tree) had the port in a Blocked state. • The source port was suppressed for multi-target packets. • The priority queue specified in the traffic counter was not allowed for some other reason. • Unknown unicast and unregistered multicast packets were filtered.
Congestion Drops	The number of outbound packets that were dropped because of traffic congestion.

Viewing egress queue counters on ICX 6610 and FCX devices

The **show interface** command displays the number of packets on a port that were queued for each QoS priority (traffic class) and dropped because of congestion.

NOTE
These counters do not include traffic on management ports or for a stack member unit that is down.

The egress queue counters display at the end of the **show interface** command output as shown in the following example.

```
Brocade#show interface e 1/1/1
GigabitEthernet1/1/1 is up, line protocol is up
  Hardware is GigabitEthernet, address is 0000.0077.8080 (bia 0000.0077.8080)
  Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
  Configured mdi mode AUTO, actual none
  Member of L2 VLAN ID 52, port is untagged, port state is FORWARDING
  BPDU guard is Disabled, ROOT protect is Disabled
  Link Error Dampening is Disabled
  STP configured to ON, priority is level0, mac-learning is enabled
  Flow Control is config enabled, oper enabled, negotiation disabled
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  Inter-Packet Gap (IPG) is 96 bit times
  IP MTU 1500 bytes
  300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  300 second output rate: 256 bits/sec, 0 packets/sec, 0.00% utilization
  0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts, 0 multicasts, 0 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 runts, 0 giants
  215704 packets output, 13805066 bytes, 0 underruns
  Transmitted 0 broadcasts, 215704 multicasts, 0 unicasts
  0 output errors, 0 collisions
  Relay Agent Information option: Disabled

Egress queues:
Queue counters    Queued packets    Dropped Packets
    0                 0                 0
    1                 0                 0
    2                 1                 0
    3                 0                 0
    4                 0                 0
    5                 0                 0
    6                 0                 0
    7               215703              0
```

Syntax: **show interface** [**ethernet** *port*]

Specify the *port* variable in the format *stack-unit/slotnum/portnum.*

Table 27 defines the egress queue statistics displayed in the output.

TABLE 27 Egress queue statistics

Parameter	Description
Queue counters	The QoS traffic class.
Queued packets	The number of packets queued on the port for the given traffic class.
Dropped packets	The number of packets for the given traffic class that were dropped because of congestion.

Clearing the egress queue counters

You can clear egress queue statistics (reset them to zero), using the **clear statistics** and **clear statistics ethernet** *port* command.

Syntax: **clear statistics** [**ethernet** *port*]

Specify the *port* variable in the format *stack-unit/slotnum/portnum.*

RMON support

The Brocade RMON agent supports the following groups. The group numbers come from the RMON specification (RFC 1757):

- Statistics (RMON Group 1)
- History (RMON Group 2)
- Alarms (RMON Group 3)
- Events (RMON Group 9)

The CLI allows you to make configuration changes to the control data for these groups, but you need a separate RMON application to view and display the data graphically.

Maximum number of entries allowed in the RMON control table

You can specify the maximum number of entries allowed in the RMON control table, including alarms, history, and events. The default number of RMON entries allowed in the RMON control table is 2048 on the FSX 800 and FSX 1600. The maximum number of RMON entries supported is 32768.

To set the maximum number of allowable entries to 3000 in the RMON history table, enter commands such as the following.

```
Brocade(config)#system-max rmon-entries 3000
Brocade(config)#write mem
Brocade(config)#exit
Brocade#reload
```

NOTE
You must save the change to the startup-config file and reload or reboot. The change does not take effect until you reload or reboot.

Syntax: **system-max rmon-entries** *value*

where *value* can be:

- 1536 – 32768 for FSX 800 and FSX 1600 devices

Statistics (RMON group 1)

Count information on multicast and broadcast packets, total packets sent, undersized and oversized packets, CRC alignment errors, jabbers, collision, fragments and dropped events is collected for each port on a Brocade Layer 2 Switch or Layer 3 Switch.

The statistics group collects statistics on promiscuous traffic across an interface. The interface group collects statistics on total traffic into and out of the agent interface.

No configuration is required to activate collection of statistics for the Layer 2 Switch or Layer 3 Switch. This activity is by default automatically activated at system start-up.

You can view a textual summary of the statistics for all ports by entering the following CLI command.

```
Brocade#show rmon statistics
Ethernet statistics 1 is active, owned by monitor
 Interface 1/1 (ifIndex 1) counters
                  Octets          0
             Drop events          0                    Packets          0
          Broadcast pkts          0             Multicast pkts          0
    CRC alignment errors          0             Undersize pkts          0
           Oversize pkts          0                  Fragments          0
                 Jabbers          0                 Collisions          0
          64 octets pkts          0      65 to 127 octets pkts          0
  128 to 255 octets pkts          0     256 to 511 octets pkts          0
 512 to 1023 octets pkts          0   1024 to 1518 octets pkts          0
```

Syntax: **show rmon statistics** [**ethernet** *port*]

NOTE
Though 48GC modules receive oversized packets and jabbers, they do not support count information for oversized packets and jabbers and the output of the **show rmon statisitics** command reports 0 for both of these counters.

The *port* parameter specifies the port number. You can use the physical port number or the SNMP port number. The physical port number is based on the product.

The SNMP numbers of the ports start at 1 and increase sequentially. For example, if you are using a Chassis device and slot 1 contains an 8-port module, the SNMP number of the first port in slot 2 is 9. The physical port number of the same port is 2/1.

This command shows the following information.

TABLE 28 Export configuration and statistics

Parameter	Definition
Octets	The total number of octets of data received on the network. This number includes octets in bad packets. This number does not include framing bits but does include Frame Check Sequence (FCS) octets.
Drop events	Indicates an overrun at the port. The port logic could not receive the traffic at full line rate and had to drop some packets as a result. The counter indicates the total number of events in which packets were dropped by the RMON probe due to lack of resources. This number is not necessarily the number of packets dropped, but is the number of times an overrun condition has been detected.
Packets	The total number of packets received. This number includes bad packets, broadcast packets, and multicast packets.
Broadcast pkts	The total number of good packets received that were directed to the broadcast address. This number does not include multicast packets.
Multicast pkts	The total number of good packets received that were directed to a multicast address. This number does not include packets directed to the broadcast address.
CRC alignment errors	The total number of packets received that were from 64 – 1518 octets long, but had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). The packet length does not include framing bits but does include FCS octets.
Undersize pkts	The total number of packets received that were less than 64 octets long and were otherwise well formed. This number does not include framing bits but does include FCS octets.
Fragments	The total number of packets received that were less than 64 octets long and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). It is normal for this counter to increment, since it counts both runts (which are normal occurrences due to collisions) and noise hits. This number does not include framing bits but does include FCS octets.
Oversize packets	The total number of packets received that were longer than 1518 octets and were otherwise well formed. This number does not include framing bits but does include FCS octets. NOTE: 48GC modules do not support count information on oversized packets and report 0.
Jabbers	The total number of packets received that were longer than 1518 octets and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). NOTE: This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms. This number does not include framing bits but does include FCS octets. NOTE: 48GC modules do not support count information on jabbers and report 0.
Collisions	The best estimate of the total number of collisions on this Ethernet segment.
64 octets pkts	The total number of packets received that were 64 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.

TABLE 28 Export configuration and statistics (Continued)

Parameter	Definition
65 to 127 octets pkts	The total number of packets received that were 65 – 127 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
128 to 255 octets pkts	The total number of packets received that were 128 – 255 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
256 to 511 octets pkts	The total number of packets received that were 256 – 511 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
512 to 1023 octets pkts	The total number of packets received that were 512 – 1023 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.
1024 to 1518 octets pkts	The total number of packets received that were 1024 – 1518 octets long. This number includes bad packets. This number does not include framing bits but does include FCS octets.

History (RMON group 2)

All active ports by default will generate two history control data entries per active Brocade Layer 2 Switch port or Layer 3 Switch interface. An active port is defined as one with a link up. If the link goes down the two entries are automatically deleted.

Two history entries are generated for each device:

- A sampling of statistics every 30 seconds
- A sampling of statistics every 30 minutes

The history data can be accessed and displayed using any of the popular RMON applications

A sample RMON history command and its syntax is shown below.

```
Brocade(config)#rmon history 1 interface 1 buckets 10 interval 10 owner nyc02
```

Syntax: **rmon history** *entry-number* **interface** *port* **buckets** *number* **interval** *sampling-interval* **owner** *text-string*

You can modify the sampling interval and the bucket (number of entries saved before overwrite) using the CLI. In the above example, owner refers to the RMON station that will request the information.

NOTE
To review the control data entry for each port or interface, enter the **show rmon history** command.

Alarm (RMON group 3)

Alarm is designed to monitor configured thresholds for any SNMP integer, time tick, gauge or counter MIB object. Using the CLI, you can define what MIB objects are monitored, the type of thresholds that are monitored (falling, rising or both), the value of those thresholds, and the sample type (absolute or delta).

An alarm event is reported each time that a threshold is exceeded. The alarm entry also indicates the action (event) to be taken if the threshold be exceeded.

A sample CLI alarm entry and its syntax is shown below.

```
Brocade(config)#rmon alarm 1 ifInOctets.6 10 delta rising-threshold 100 1 falling
threshold 50 1 owner nyc02
```

Syntax: **rmon alarm** *entry-number MIB-object.interface-num sampling-time sample-type threshold-type threshold-value event-number threshold-type threshold-value event-number*
owner *text-string*

Event (RMON group 9)

There are two elements to the Event Group—the ***event control table*** and the ***event log table***.

The event control table defines the action to be taken when an alarm is reported. Defined events can be found by entering the CLI command, show event. The Event Log Table collects and stores reported events for retrieval by an RMON application.

A sample entry and syntax of the event control table is shown below.

```
Brocade(config)#rmon event 1 description 'testing a longer string' log-and-trap
public owner nyc02
```

Syntax: **rmon event** *event-entry* **description** *text-string* **log** | **trap** | **log-and-trap owner** *rmon-station*

sFlow

NOTE
FastIron devices support sFlow version 5 by default.

sFlow is a standards-based protocol that allows network traffic to be sampled at a user-defined rate for the purpose of monitoring traffic flow patterns and identifying packet transfer rates on user-specified interfaces.

When sFlow is enabled on a Layer 2 or Layer 3 switch, the system performs the following sFlow-related tasks:

- Samples traffic flows by copying packet header information
- Identifies ingress and egress interfaces for the sampled flows
- Combines sFlow samples into UDP packets and forwards them to the sFlow collectors for analysis
- Forwards byte and packet count data, or counter samples, to sFlow collectors

sFlow is described in RFC 3176, “InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks”.

On ICX and FCX Series devices, you can use QoS queue 1 for priority traffic, even when sFlow is enabled on the port. This differs from FastIron X Series devices, which support seven priorities instead of eight when sFlow is enabled. In this case, QoS queue 1 is reserved for sFlow and is not used by other packets. Any non-sFlow packets assigned to QoS queue 1 will be directed to QoS queue 0.

sFlow version 5

sFlow version 5 enhances and modifies the format of the data sent to the sFlow collector. sFlow version 5 introduces several new sFlow features and also defines a new datagram syntax used by the sFlow agent to report flow samples and interface counters to the sFlow collector.

sFlow version 5 adds support for the following:

- sFlow version 5 datagrams
- Sub-agent support
- Configurable sFlow export packet size
- Support for the new data field and sample type length in flow samples
- Configurable interval for exporting Brocade-specific data structure

sFlow version 5 is backward-compatible with sFlow version 2. By default, the sFlow agent exports sFlow version 5 flow samples by default, but you can configure the device to export the data in sFlow version 2 format. You can switch between sFlow version 2 and sFlow version 5 formats. The sFlow collector automatically parses each incoming sample and decodes it based on the version number.

The configuration procedures for sFlow version 5 are the same as for sFlow version 2, except where explicitly noted. Configuration procedures for sFlow are in the section "Configuring and enabling sFlow" on page 275. The features and CLI commands that are specific to sFlow version 5 are described in the section "sFlow version 5 feature configuration" on page 282.

sFlow support for IPv6 packets

The Brocade implementation of sFlow features support IPv6 packets. This support includes extended router information and extended gateway information in the sampled packet. Note that sFlow support for IPv6 packets exists only on devices running software that supports IPv6.

The configuration procedures for this feature are the same as for IPv4, except where the collector is a link-local address on a Layer 3 switch. For details refer to "Specifying the collector" on page 275.

Extended router information

IPv6 sFlow sampled packets include the following extended router information:

- IP address of the next hop router
- Outgoing VLAN ID
- Source IP address prefix length
- Destination IP address prefix length

Note that in IPv6 devices, the prefix lengths of the source and destination IP addresses are collected if BGP is configured and the route lookup is completed. In IPv4 devices, this information is collected only if BGP is configured on the devices.

Extended gateway information

If BGP is enabled, extended gateway information is included in IPv6 sFlow sampled packets, including the following BGP information about a packet destination route:

- The Autonomous System number for the router

- The source IP Autonomous System of the route
- The source peer Autonomous System for the route
- The Autonomous System patch to the destination

NOTE
Autonomous System communities and local preferences are not included in the sampled packets.

To obtain extended gateway information, use “struct extended_gateway” as described in RFC 3176.

IPv6 packet sampling

IPv6 sampling is performed by the packet processor. The system uses the sampling rate setting to selectively mark the monitoring bit in the header of an incoming packet. Marked packets tell the CPU that the packets are subject to sFlow sampling.

sFlow configuration considerations

This section lists the sFlow configuration considerations on Brocade devices.

On ICX and FCX Series devices, you can use QoS queue 1 for priority traffic, even when sFlow is enabled on the port. This differs from FastIron X Series devices, which support seven priorities instead of eight when sFlow is enabled. In this case, QoS queue 1 is reserved for sFlow and is not used by other packets. Any non-sFlow packets assigned to QoS queue 1 will be directed to QoS queue 0.

If ICX and FCX stacks are rebooted, sFlow is disabled on standby and member units until the configuration is synchronized between the Active and Standby Controllers.

sFlow and hardware support

- Brocade devices support sFlow packet sampling of inbound traffic only. These devices do not sample outbound packets. However, Brocade devices support byte and packet count statistics for both traffic directions.
- sFlow is supported on all Ethernet ports (10/100, Gbps, and 10 Gbps)

sFlow and CPU utilization

Enabling sFlow may cause a slight and noticeable increase of up to 20% in CPU utilization. In typical scenarios, this is normal behavior for sFlow, and does not affect the functionality of other features on the switch.

sFlow and source address

The sampled sFlow data sent to the collectors includes an agent_address field. This field identifies the IP address of the device that sent the data:

- On a Layer 2 Switch, agent_address is the Layer 2 Switch management IP address. You must configure the management IP address in order to export sFlow data from the device. If the switch has both an IPv4 and IPv6 address, the agent_address is the IPv4 address. If the switch has an IPv6 address only, the agent_address is the global IPv6 address.

- On a Layer 3 Switch with IPv6 interfaces only, sFlow looks for an IPv6 address in the following order, and uses the first address found:
  - The first IPv6 address on the lowest-numbered loopback interface
  - The first IPv6 address on the lowest-numbered VE interface
  - The first IPv6 address on any interface
- On a Layer 3 Switch with both IPv4 and IPv6 interfaces, or with IPv4 interfaces only, sFlow looks for an IP address in the following order, and uses the first address found:
  - The IPv4 router ID configured by the **ip router-id** command
  - The first IPv4 address on the lowest-numbered loopback interface
  - The first IPv4 address on the lowest-numbered virtual interface
  - The first IPv4 address on any interface

NOTE
The device uses the router ID only if the device also has an IP interface with the same address. Router ID is not supported on IPv6 devices.

NOTE
If an IP address is not already configured when you enable sFlow, the feature uses the source address 0.0.0.0. To display the agent_address, enable sFlow, then enter the **show sflow** command. Refer to "Enabling sFlow forwarding" on page 281 and "Displaying sFlow information" on page 285.

NOTE
In sFlow version 5, you can set an arbitrary IPv4 or IPv6 address as the sFlow agent IP address. Refer to "Specifying the sFlow agent IP address" on page 283.

sFlow and source port

By default, sFlow sends data to the collector out of UDP source port 8888, but you can specify a different source port. For more information, refer to "Changing the sFlow source port" on page 281.

sFlow and sampling rate

The ***sampling rate*** is the average ratio of the number of packets incoming on an sFlow enabled port, to the number of flow samples taken from those packets. sFlow sampling can affect performance in some configurations.

Note that on the FastIron devices, the configured sampling rate and the actual rate are the same. The software does not adjust the configured sampling rate as on other Brocade devices.

sFlow and port monitoring

- ICX and FCX Series devices support sFlow and port monitoring together on the same port.

- FastIron X Series devices support port monitoring and sFlow together on the same device. The caveat is that these features cannot be configured together within the same port region on non-third generation modules. The following third-generation SX modules support sFlow and mirroring on the same port:
  - SX-FI48GPP
  - SX-FI-24GPP
  - SX-FI-24HF
  - SX-FI-2XG
  - SX-FI-8XG

Configuring and enabling sFlow

NOTE
The commands in this section apply to sFlow version 2 and sFlow version 5. CLI commands that are specific to sFlow version 5 are documented in "sFlow version 5 feature configuration" on page 282.

To configure sFlow, perform the following tasks:

- Optional – If your device supports sFlow version 5, change the version used for exporting sFlow data
- Specify collector information. The collector is the external device to which you are exporting the sFlow data. You can specify up to four collectors.
- Optional – Change the polling interval
- Optional – Change the sampling rate
- Optional – Change the sFlow source port
- Enable sFlow globally
- Enable sFlow forwarding on individual interfaces
- Enable sFlow forwarding on individual trunk ports
- If your device supports sFlow version 5, configure sFlow version 5 features

NOTE
If you change the router ID or other IP address value that sFlow uses for its agent_address, you need to disable and then re-enable sFlow to cause the feature to use the new source address.

Specifying the collector

sFlow exports traffic statistics to an external collector. You can specify up to four collectors. You can specify more than one collector with the same IP address if the UDP port numbers are unique. You can have up to four unique combinations of IP addresses and UDP port numbers.

Specifying an sFlow collector on IPv4 devices

To specify an sFlow collector on an IPv4 device, enter a command such as the following.

```
Brocade(config)#sflow destination 10.10.10.1
```

This command specifies a collector with IPv4 address 10.10.10.1, listening for sFlow data on UDP port 6343.

Syntax: [**no**] **sflow destination** *ip-addr* [*dest-udp-port* | *vrf*]

The *ip-addr* parameter specifies the IP address of the collector.

The *dest-udp-port* parameter specifies the UDP port on which the sFlow collector will be listening for exported sFlow data. The default port number is 6343. For information on VRF parameter, see the *FastIron Layer 3 Routing Configuration Guide*.

The sampled sFlow data sent to the collectors includes an agent_address field. This field identifies the device that sent the data. Refer to "sFlow and source address" on page 273.

Specifying an sFlow collector on IPv6 devices

To specify an sFlow collector on an IPv6 device, enter a command such as the following.

```
Brocade(config)#sflow destination ipv6 2001:DB8:0::0b:02a
```

This command specifies a collector with IPv6 address 2001:DB8::0b:02a, listening for sFlow data on UDP port 6343.

Syntax: [**no**] **sflow destination ipv6** *ip-addr* [*dest-udp-port*]

The *ip-addr* parameter specifies the IP address of the collector.

The *dest-udp-port* parameter specifies the UDP port on which the sFlow collector will be listening for exported sFlow data. The default port number is 6343.

If the IPv6 address you specify is a link-local address on a Layer 3 switch, you must also specify the **outgoing-interface ethernet** *port-num* or the **ve** *port-num*. This identifies the outgoing interface through which the sampled packets will be sent.

The sampled sFlow data sent to the collectors includes an agent_address field. This field identifies the device that sent the data. Refer to “sFlow and source address” on page 273.

Changing the polling interval

The polling interval defines how often sFlow byte and packet counter data for a port are sent to the sFlow collectors. If multiple ports are enabled for sFlow, the Brocade device staggers transmission of the counter data to smooth performance. For example, if sFlow is enabled on two ports and the polling interval is 20 seconds, the Brocade device sends counter data every ten seconds. The counter data for one of the ports are sent after ten seconds, and counter data for the other port are sent after an additional ten seconds. Ten seconds later, new counter data for the first port are sent. Similarly, if sFlow is enabled on five ports and the polling interval is 20 seconds, the Brocade device sends counter data every four seconds.

The default polling interval is 20 seconds. You can change the interval to a value from 1 to any higher value. The interval value applies to all interfaces on which sFlow is enabled. If you set the polling interval to 0, counter data sampling is disabled.

To change the polling interval, enter a command such as the following at the global CONFIG level of the CLI.

```
Brocade(config)#sflow polling-interval 30
```

Syntax: [**no**] **sflow polling-interval** *secs*

The *secs* parameter specifies the interval and can be from 1 to any higher value. The default is 20 seconds. If you specify 0, counter data sampling is disabled.

Changing the sampling rate

The sampling rate is the average ratio of the number of packets incoming on an sFlow-enabled port, to the number of flow samples taken from those packets.

You can change the default (global) sampling rate. You also can change the rate on an individual port, overriding the default sampling rate of 512. With a sampling rate of 512, on average, one in every 512 packets forwarded on an interface is sampled.

Configuration considerations

The sampling rate is a fraction in the form 1/N, meaning that, on average, one out of every N packets will be sampled. The **sflow sample** command at the global level or port level specifies N, the denominator of the fraction. Thus a higher number for the denominator means a lower sampling rate since fewer packets are sampled. Likewise, a lower number for the denominator means a higher sampling rate because more packets are sampled. For example, if you change the denominator from 512 to 128, the sampling rate increases because four times as many packets will be sampled.

NOTE
Brocade recommends that you do not change the denominator to a value lower than the default. Sampling requires CPU resources. Using a low denominator for the sampling rate can cause high CPU utilization.

Configured rate and actual rate

When you enter a sampling rate value, this value is the *configured rate* as well as the *actual sampling rate*.

Change to global rate

If you change the global sampling rate, the change is applied to all sFlow-enabled ports *except* those ports on which you have already explicitly set the sampling rate. For example, suppose that sFlow is enabled on ports 1/1, 1/2, and 5/1. If you configure the sampling rate on port 1/1 but leave the other two ports using the default rate, then a change to the global sampling rate applies to ports 1/2 and 5/1 but not port 1/1. sFlow assumes that you want to continue using the sampling rate you explicitly configured on an individual port even if you globally change the sampling rate for the other ports.

Module rate

While different ports on a module may be configured to have different sampling rates, the hardware for the module will be programmed to take samples at a single rate (the module sampling rate). The module sampling rate will be the highest sampling rate (i.e. lowest number) configured for any of the ports on the module.

When ports on a given module are configured with different sampling rates, the CPU discards some of the samples supplied by the hardware for ports with configured sampling rates which are lower than the module sampling rate. This is referred to as subsampling, and the ratio between the port sampling rate and the module sampling rate is known as the subsampling factor. For example, if the module in slot 4 has sFlow enabled on ports 4/2 and 4/8, and port 4/2 is using the default sampling rate of 512, and port 4/8 is configured explicitly for a rate of 2048, then the module sampling rate will be 512 because this is this highest port sampling rate (lowest number). The subsampling factor for port 4/2 will be 1, meaning that every sample taken by the hardware will be exported, while the subsampling factor for port 4/8 will be 4, meaning that one out of every four samples taken by the hardware will be exported. Whether a port's sampling rate is configured explicitly, or whether it uses the global default setting, has no effect on the calculations.

You do not need to perform any of these calculations to change a sampling rate. For simplicity, the syntax information in this section lists the valid sampling rates. You can display the rates you entered for the default sampling rate, module rates, and all sFlow-enabled ports by entering the **show sflow** command. Refer to "Displaying sFlow information" on page 285.

Sampling rate for new ports

When you enable sFlow on a port, the port's sampling rate is set to the global default sampling rate. This also applies to ports on which you disable and then re-enable sFlow. The port does not retain the sampling rate it had when you disabled sFlow on the port, even if you had explicitly set the sampling rate on the port.

Changing the default sampling rate

To change the default (global) sampling rate, enter a command such as the following at the global CONFIG level of the CLI.

```
Brocade(config)#sflow sample 2048
```

Syntax: [**no**] **sflow sample** *num*

The *num* parameter specifies the average number of packets from which each sample will be taken. The software rounds the value you enter to the next higher odd power of 2. This value becomes the actual default sampling rate and is one of the following:

- 2
- 8
- 32
- 128
- 512
- 2048
- 4096
- 8192
- 32768
- 131072
- 524288
- 2097152
- 8388608
- 33554432
- 134217728
- 536870912
- 2147483648

For example, if the configured sampling rate is 1000, then the actual rate is 2048 and 1 in 2048 packets are sampled by the hardware.

Changing the sampling rate of a module

You cannot change a module sampling rate directly. You can change a module sampling rate only by changing the sampling rate of a port on that module.

Changing the sampling rate on a port

You can configure an individual port to use a different sampling rate than the global default sampling rate. This is useful in cases where ports have different bandwidths. For example, if you are using sFlow on 10/100 ports and Gbps Ethernet ports, you might want to configure the Gbps ports to use a higher sampling rate (and thus gather fewer samples per number of packets) than the 10/100 ports.

To change the sampling rate on an individual port, enter a command such as the following at the configuration level for the port.

```
Brocade(config-if-1/1)#sflow sample 8192
```

Syntax: [**no**] **sflow sample** *num*

The *num* parameter specifies the average number of packets from which each sample will be taken. The software rounds the value you enter up to the next odd power of 2. The actual sampling rate becomes one of the values listed in "Changing the default sampling rate".

NOTE
Configuring a sampling rate on a port that is the primary port of a trunk applies that same sampling rate to all ports in the trunk.

Changing the sampling rate for a trunk port

You can configure an individual static trunk port to use a different sampling rate than the global default sampling rate. This feature is also supported on LACP trunk ports. This feature is useful in cases where ports have different bandwidths. For example, if you are using sFlow on 10/100 ports and Gbps Ethernet ports, you might want to configure the Gbps ports to use a higher sampling rate (and thus gather fewer samples per number of packets) than the 10/100 ports.

To configure a static trunk port to use a different sampling rate than the global default sampling rate, enter commands such as the following:

```
Brocade(config)#trunk e 4/1 to 4/8
Brocade(config-trunk-4/1-4/8)sflow sample 8192
```

Syntax: [**no**] **sflow sample** *num*

The *num* parameter specifies the average number of packets from which each sample will be taken. The software rounds the value you enter up to the next odd power of 2. The actual sampling rate becomes one of the values listed in "Changing the default sampling rate".

NOTE
Configuring a sampling rate on only the port that is the primary port of a trunk automatically applies that same sampling rate to all ports in the trunk.

Changing the sFlow source port

By default, sFlow sends data to the collector using UDP source port 8888, but you can change the source UDP port to any port number in the range 1025-65535.

To change the source UDP port, enter a command such as the following:

```
Brocade(config)#sflow source-port 8000
```

Syntax: [no] **sflow source-port** *num*

The *num* parameter specifies the sFlow source port.

Enabling sFlow forwarding

sFlow exports data only for the interfaces on which you enable sFlow forwarding. You can enable sFlow forwarding on Ethernet interfaces.

To enable sFlow forwarding, perform the following:

- Globally enable the sFlow feature
- Enable sFlow forwarding on individual interfaces
- Enable sFlow forwarding on individual trunk ports

NOTE
Before you enable sFlow, make sure the device has an IP address that sFlow can use as its source address. Refer to "sFlow and source address" on page 273 for the source address requirements.

NOTE
When you enable sFlow forwarding on an 802.1X-enabled interface, the samples taken from the interface include the username used to obtain access to either or both the inbound and outbound ports, if that information is available. For information about 802.1X, refer to "802.1X Port Security" chapter in the *FastIron Ethernet Switch Security Configuration Guide.*

Command syntax for enabling sFlow forwarding

This section shows how to enable sFlow forwarding.

Globally enabling sFlow forwarding

To enable sFlow forwarding, you must first enable it on a global basis, then on individual interfaces or trunk ports, or both.

To globally enable sFlow forwarding, enter the following command.

```
Brocade(config)#sflow enable
```

You can now enable sFlow forwarding on individual ports as described in the next two sections.

Syntax: [no] **sflow enable**

Enabling sFlow forwarding on individual interfaces

To enable sFlow forwarding enter commands such as the following.

```
Brocade(config)#sflow enable
Brocade(config)#interface ethernet 1/1 to 1/8
Brocade(config-mif-1/1-1/8)#sflow forwarding
```

These commands globally enable sFlow, then enable sFlow forwarding on Ethernet ports 1/1 – 1/8. You must use both the **sflow enable** and **sflow forwarding** commands to enable the feature.

Syntax: [**no**] **sflow enable**

Syntax: [**no**] **sflow forwarding**

Enabling sFlow forwarding on individual trunk ports

This feature is supported on individual ports of a static trunk group. It is also supported on LACP trunk ports.

NOTE
When you enable sFlow forwarding on a trunk port, only the primary port of the trunk group forwards sFlow samples.

To enable sFlow forwarding on a trunk port, enter commands such as the following.

```
Brocade(config)#sflow enable
Brocade(config)#trunk e 4/1 to 4/8
Brocade(config-trunk-4/1-4/8)#config-trunk-ind
Brocade(config-trunk-4/1-4/8)#sflow forwarding e 4/2
```

These commands globally enable sFlow, then enable sFlow forwarding on trunk port e 4/2. You must use both the **sflow enable** and **sflow forwarding** commands to enable the feature.

Syntax: [**no**] **sflow enable**

Syntax: [**no**] **sflow forwarding**

sFlow version 5 feature configuration

NOTE
The commands in this section are supported when sFlow version 5 is enabled on the device. These commands are not supported with sFlow version 2. sFlow version 5 also supports all of the sFlow configuration commands in "Configuring and enabling sFlow" on page 275.

When sFlow version 5 is enabled on the device, you can do the following:

- Specify the sFlow version (version 2 or version 5)
- Specify the sFlow agent IP address
- Specify the maximum flow sample size
- Export CPU and memory usage Information to the sFlow collector
- Specify the polling interval for exporting CPU and memory usage information to the sFlow collector
- Export CPU-directed data (management traffic) to the sFlow collector

Egress interface ID for sampled broadcast and multicast packets

For broadcast and multicast traffic, the egress interface ID for sampled traffic is always 0x80000000. When broadcast and multicast packets are sampled, they are usually forwarded to more than one port. However, the output port field in an sFlow datagram supports the display of one egress interface ID only. Therefore, the sFlow version 5 agent always sets the output port ID to 0x80000000 for broadcast and multicast packets that are sampled.

Specifying the sFlow version format

If your device supports sFlow version 5, you can optionally specify the version used for exporting sFlow data. Refer "Specifying the sFlow agent IP address".

Specifying the sFlow agent IP address

The sampled sFlow data sent to the collectors includes an agent_address field. This field identifies the device (the sFlow agent) that sent the data. By default, the device automatically selects the sFlow agent IP address based on the configuration, as described in the section "sFlow and source address" on page 273. Alternatively, you can configure the device to instead use an arbitrary IPv4 or IPv6 address as the sFlow agent IP address.

To specify an IPv4 address as the sFlow agent IP address, enter a command such as the following

```
Brocade(config)#sflow agent-ip 10.10.10.1
```

Syntax: [**no**] **sflow agent-ip** *ipv4-addr*

The *ipv4-addr* specifies the address of the device that sent the data.

To specify an IPv6 address as the sFlow agent IP address, enter a command such as the following.

```
Brocade(config)#sflow agent-ip FE80::240:D0FF:FE48:4672
```

Syntax: [**no**] **sflow agent-ip** *ipv6-addr*

The *ipv6-addr* the address of the device that sent the data.

Specifying the version used for exporting sFlow data

By default, when sFlow is enabled globally on the Brocade device, the sFlow agent exports sFlow data in version 5 format. You can change this setting so that the sFlow agent exports data in version 2 format. You can switch between versions without rebooting the device or disabling sFlow.

NOTE
When the sFlow version number is changed, the system will reset sFlow counters and flow sample sequence numbers.

To specify the sFlow version used for exporting sFlow data, enter the following command.

```
Brocade(config)#sflow version 2
```

Syntax: [**no**] **sflow version** 2 | 5

The default is 5.

Specifying the maximum flow sample size

With sFlow version 5, you can specify the maximum size of the flow sample sent to the sFlow collector. If a packet is larger than the specified maximum size, then only the contents of the packet up to the specified maximum number of bytes is exported. If the size of the packet is smaller than the specified maximum, then the entire packet is exported.

For example, to specify 1024 bytes as the maximum flow sample size, enter the following command.

```
Brocade(config)# sflow max-packet-size 1024
```

Syntax: [**no**] **sflow max-packet-size** *size*

For both sFlow version 2 and version 5, the default maximum flow sample size is 256 bytes.

For sFlow version 5, the maximum flow sample size is 1300 bytes.

Exporting CPU and memory usage information to the sFlow collector

With sFlow version 5, you can optionally configure the sFlow agent on the Brocade device to export information about CPU and memory usage to the sFlow collector.

To export CPU usage and memory usage information, enter the following command.

```
Brocade(config)# sflow export system-info
```

Syntax: [**no**] **sflow export system-info**

By default, CPU usage information and memory usage information are not exported.

Specifying the polling interval for exporting CPU and memory usage information to the sFlow collector

The polling interval defines how often sFlow data for a port is sent to the sFlow collector. With sFlow version 5, you can optionally set the polling interval used for exporting CPU and memory usage information.

For example, to set the polling interval for exporting CPU and memory usage information to 30 seconds, enter the following command.

```
Brocade(config)# sflow export system-info 30
```

Syntax: [**no**] **sflow export system-info** *seconds*

You can specify a polling interval from 5 seconds to 1,800 seconds (30 minutes). The default polling interval for exporting CPU and memory usage information is 300 seconds (5 minutes).

Exporting CPU-directed data (management traffic) to the sFlow collector

You can select which and how often data destined to the CPU (for example, Telnet sessions) is sent to the sFlow collector.

CLI commands allow you to do the following:

- Enable the sFlow agent to export CPU-directed data
- Specify the sampling rate for exported CPU-directed data

Enabling the sFlow agent to export CPU-directed data

To enable the sFlow agent on a Brocade device to export data destined to the CPU to the sFlow collector, enter the following command.

```
Brocade(config)# sflow export cpu-traffic
```

Syntax: [**no**] **sflow export cpu-traffic**

By default, this feature is disabled. The sFlow agent does not send data destined to the CPU to the sFlow collector.

Specifying the sampling rate for exported CPU-directed data

The sampling rate is the average ratio of the number of packets incoming on an sFlow-enabled port, to the number of flow samples taken from those packets. You can optionally set the sampling rate for CPU-directed data exported to the sFlow collector. For example, to set this sampling rate to 2048, enter the following command.

```
Brocade(config)# sflow export cpu-traffic 2048
```

Syntax: [**no**] **sflow export cpu-traffic** *rate*

The default sampling rate depends on the Brocade device being configured. Refer to "Changing the sampling rate" on page 277 for the default sampling rate for each kind of Brocade device.

Displaying sFlow information

To display sFlow configuration information and statistics, enter the following command at any level of the CLI.

```
Brocade#show sflow
sFlow version:5
sFlow services are enabled.
sFlow agent IP address: 10.123.123.1
4 collector destinations configured:
Collector IP 192.168.4.204, UDP 6343
Collector IP 192.168.4.200, UDP 6333
Collector IP 192.168.4.202, UDP 6355
Collector IP 192.168.4.203, UDP 6565
Configured UDP source port: 33333
Polling interval is 0 seconds.
Configured default sampling rate: 1 per 512 packets
Actual default sampling rate: 1 per 512 packets
The maximum sFlow sample size:512
exporting cpu-traffic is enabled
exporting cpu-traffic sample rate:16
exporting system-info is enabled
exporting system-info polling interval:20 seconds
10552 UDP packets exported
24127 sFlow samples collected.
sFlow ports: ethe 1/2 to 1/12 ethe 1/15 ethe 1/25 to 1/26 ethe 4/1 ethe 5/10 to
5/20 ethe 8/1 ethe 8/4
Module Sampling Rates
---------------------
Slot  1 configured rate=512, actual rate=512
Slot  3 configured rate=0, actual rate=0
Slot  4 configured rate=10000, actual rate=32768
Slot  5 configured rate=512, actual rate=512
Slot  7 configured rate=0, actual rate=0
Slot  8 configured rate=512, actual rate=512
Port Sampling Rates
-------------------
Port 8/4, configured rate=512, actual rate=512, Subsampling factor=1
Port 8/1, configured rate=512, actual rate=512, Subsampling factor=1
Port 5/20, configured rate=3000, actual rate=8192, Subsampling factor=16
Port 5/19, configured rate=512, actual rate=512, Subsampling factor=1
Port 5/18, configured rate=512, actual rate=512, Subsampling factor=1
Port 5/17, configured rate=1500, actual rate=2048, Subsampling factor=4
Port 5/16, configured rate=1500, actual rate=2048, Subsampling factor=4
Port 5/15, configured rate=1500, actual rate=2048, Subsampling factor=4
Port 5/14, configured rate=1500, actual rate=2048, Subsampling factor=4
Port 5/13, configured rate=512, actual rate=512, Subsampling factor=1
Port 5/12, configured rate=512, actual rate=512, Subsampling factor=1
Port 5/11, configured rate=512, actual rate=512, Subsampling factor=1
Port 5/10, configured rate=512, actual rate=512, Subsampling factor=1
Port 4/1, configured rate=10000, actual rate=32768, Subsampling factor=1
Port 1/26, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/25, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/15, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/12, configured rate=512, actual rate=512, Subsampling factor=1
```

...continued on next page...

```
...continued from previous page...
Port 1/11, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/10, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/9, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/8, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/7, configured rate=1000, actual rate=2048, Subsampling factor=4
Port 1/6, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/5, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/4, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/3, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/2, configured rate=1000, actual rate=2048, Subsampling factor=4
```

Syntax: **show sflow**

The show sflow command displays the following information.

TABLE 29 sFlow information

Parameter	Definition
sFlow version	The version of sFlow enabled on the device, which can be one of the following: • 2 • 5
sFlow services	The feature state, which can be one of the following: • disabled • enabled
sFlow agent IP address	The IP address that sFlow is using in the agent_address field of packets sent to the collectors. Refer to "sFlow and source address" on page 273.
Collector	The collector information. The following information is displayed for each collector: • IP address • UDP port If more than one collector is configured, the line above the collectors indicates how many have been configured.
Configured UDP source port	The UDP source port used to send data to the collector.
Polling interval	The port counter polling interval.
Configured default sampling rate	The configured global sampling rate. If you changed the global sampling rate, the value you entered is shown here. The actual rate calculated by the software based on the value you entered is listed on the next line, "Actual default sampling rate".
Actual default sampling rate	The actual default sampling rate.
The maximum sFlow sample size	The maximum size of a flow sample sent to the sFlow collector.
exporting cpu-traffic	Indicates whether or not the sFlow agent is configured to export data destined to the CPU (e.g., Telnet sessions) to the sFlow collector: • enabled • disabled
exporting cpu-traffic sample rate	The sampling rate for CPU-directed data, which is the average ratio of the number of incoming packets on an sFlow-enabled port, to the number of flow samples taken from those packets.

TABLE 29 sFlow information (Continued)

Parameter	Definition
exporting system-info	Indicates whether or not the sFlow agent is configured to export information about CPU and memory usage to the sFlow collector: • enabled • disabled
exporting system-info polling interval	Specifies the interval, in seconds, that sFlow data is sent to the sFlow collector.
UDP packets exported	The number of sFlow export packets the Brocade device has sent. **NOTE:** Each UDP packet can contain multiple samples.
sFlow samples collected	The number of sampled packets that have been sent to the collectors.
sFlow ports	The ports on which you enabled sFlow.
Module Sampling Rates	The configured and actual sampling rates for each module. If a module does not have any sFlow-enabled ports, the rates are listed as 0.
Port Sampling Rates	The configured and actual sampling rates for each sFlow-enabled port. The Subsampling factor indicates how many times the sampling rate of the port's module is multiplied to achieve the port's sampling rate. Because of the way the actual sampling rates are computed, the Subsampling factors are always whole numbers.

Clearing sFlow statistics

To clear the UDP packet and sFlow sample counters in the **show sflow** display, enter the following command.

```
Brocade#clear statistics
```

Syntax: **clear statistics**

This command clears the values in the following fields of the **show sflow** display:

- UDP packets exported
- sFlow samples collected

NOTE
This command also clears the statistics counters used by other features.

Utilization list for an uplink port

You can configure uplink utilization lists that display the percentage of a given uplink port bandwidth that is used by a specific list of downlink ports. The percentages are based on 30-second intervals of RMON packet statistics for the ports. Both transmit and receive traffic is counted in each percentage.

NOTE
This feature is intended for ISP or collocation environments in which downlink ports are dedicated to various customers' traffic and are isolated from one another. If traffic regularly passes between the downlink ports, the information displayed by the utilization lists does not provide a clear depiction of traffic exchanged by the downlink ports and the uplink port.

Each uplink utilization list consists of the following:

- Utilization list number (1, 2, 3, or 4)
- One or more uplink ports
- One or more downlink ports

Each list displays the uplink port and the percentage of that port bandwidth that was utilized by the downlink ports over the most recent 30-second interval.

You can configure up to four bandwidth utilization lists.

Utilization list for an uplink port command syntax

To configure an uplink utilization list, enter commands such as the following. The commands in this example configure a link utilization list with port 1/1 as the uplink port and ports 1/2 and 1/3 as the downlink ports.

```
Brocade(config)#relative-utilization 1 uplink eth 1/1 downlink eth 1/2 to 1/3
Brocade(config)#write memory
```

Syntax: [**no**] **relative-utilization** *num* **uplink ethernet** *port* [**to** *port* | *port*...] **downlink ethernet** *port* [**to** *port* | [*port*...]

The *num* parameter specifies the list number. You can configure up to four lists. Specify a number from 1 – 4.

The **uplink ethernet** parameters and the port numbers you specify after the parameters indicate the uplink ports.

The **downlink ethernet** parameters and the port numbers you specify after the parameters indicate the downlink ports.

Displaying utilization percentages for an uplink

After you configure an uplink utilization list, you can display the list to observe the percentage of the uplink bandwidth that each of the downlink ports used during the most recent 30-second port statistics interval. The number of packets sent and received between the two ports is listed, as well as the ratio of each individual downlink port packets relative to the total number of packets on the uplink.

To display an uplink utilization list, enter a command such as the following at any level of the CLI.

```
Brocade#show relative-utilization 1
uplink: ethe 1
30-sec total uplink packet count = 3011
packet count ratio (%)
  1/ 2:60   1/ 3:40
```

In this example, ports 1/2 and 1/3 are sending traffic to port 1/1. Port 1/2 and port 1/3 are isolated (not shared by multiple clients) and typically do not exchange traffic with other ports except for the uplink port, 1/1.

Syntax: **show relative-utilization** *num*

The *num* parameter specifies the list number.

NOTE
The example above represents a pure configuration in which traffic is exchanged only by ports 1/2 and 1/1, and by ports 1/3 and 1/1. For this reason, the percentages for the two downlink ports equal 100%. In some cases, the percentages do not always equal 100%. This is true in cases where the ports exchange some traffic with other ports in the system or when the downlink ports are configured together in a port-based VLAN.

In the following example, ports 1/2 and 1/3 are in the same port-based VLAN.

```
Brocade#show relative-utilization 1
uplink: ethe 1
30-sec total uplink packet count = 3011
packet count ratio (%)
  1/ 2:100   1/ 3:100
```

Here is another example showing different data for the same link utilization list. In this example, port 1/2 is connected to a hub and is sending traffic to port 1/1. Port 1/3 is unconnected.

```
Brocade#show relative-utilization 1
uplink: ethe 1
30-sec total uplink packet count = 2996
packet count ratio (%)
  1 /2:100   1/ 3:---
```

Chapter

Power over Ethernet 12

Table 30 lists the individual Brocade FastIron switches and the Power over Ethernet (PoE) features they support. These features are supported in the Layer 2 and Layer 3 software images, except where noted.

TABLE 30 Supported PoE features

Feature	FSX 800 FSX 1600 PoE interface modules[1]	FCX PoE+ models only	ICX 6610 PoE models only	ICX 6430 PoE models only	ICX 6450 PoE models only
PoE+ (802.3at)	Yes (FSX 800 and FSX 1600 with SX-FI48GPP module only)	Yes	Yes	Yes	Yes
PoE (802.3af)	Yes	Yes	Yes	Yes	Yes
Detection of PoE power requirements advertised through CDP	Yes	Yes	Yes	Yes	Yes
Maximum power level for a PoE power consuming device	Yes	Yes	Yes	Yes	Yes
Power class for PoE power consuming device	Yes	Yes	Yes	Yes	Yes
Maximum power budget per PoE interface module	Yes	No	No	Yes	Yes
In-line power priority for a PoE port	Yes	Yes	Yes	Yes	Yes
PoE firmware upgrade via CLI	Yes	Yes	Yes	Yes	Yes

1. Supported on PoE-enabled interface modules installed in a chassis with PoE power supply.

This chapter contains the following sections:

Power over Ethernet overview

This section provides an overview of the requirements for delivering power over the LAN, as defined by the Institute of Electrical and Electronics Engineers Inc. (IEEE) in the 802.3af (PoE) and 802.3at (PoE+) specifications.

Brocade PoE devices provide Power over Ethernet, compliant with the standards described in the IEEE 802.3af specification for delivering inline power. Brocade PoE+ devices are compliant with both the 802.3af and 802.3at specifications. The 802.3af specification defined the original standard for delivering power over existing network cabling infrastructure, enabling multicast-enabled full streaming audio and video applications for converged services, such as, Voice over IP (VoIP), Wireless Local Area Access (WLAN) points, IP surveillance cameras, and other IP technology devices. The 802.3at specification expands the standards to support higher power levels for more demanding powered devices, such as video IP phones, pan-tilt-zoom cameras and high-power outdoor antennas for wireless access points. Except where noted, this document will use the term PoE to refer to both PoE and PoE+.

Table 30 lists the FastIron devices and modules that support PoE, PoE+, or both.

PoE technology eliminates the need for an electrical outlet and dedicated UPS near IP powered devices. With power sourcing equipment such as a Brocade FastIron PoE device, power is consolidated and centralized in the wiring closets, improving the reliability and resiliency of the network. Because PoE can provide Power over Ethernet cable, power is continuous, even in the event of a power failure.

Power over Ethernet terms used in this chapter

The following terms are introduced in this chapter:

- Power-sourcing device or Power-sourcing equipment (PSE) - This is the source of the power, or the device that integrates the power onto the network. Power sourcing devices and equipment have embedded PoE technology. The Brocade FastIron PoE device is a power sourcing device.
- IP powered device (PD) or power-consuming device - This is the Ethernet device that requires power and is situated on the end of the cable opposite the power sourcing equipment.

Methods for delivering Power over Ethernet

There are two methods for delivering Power over Ethernet (PoE), as defined in the 802.3af and 802.3at specifications:

- Endspan - Power is supplied through the Ethernet ports on a power sourcing device. With the Endspan solution, power can be carried over the two data pairs (Alternative A) or the two spare pairs (Alternative B).
- Midspan - Power is supplied by an intermediate power sourcing device placed between the switch and the PD. With the Midspan solution, power is carried over the two spare pairs (Alternative B).

With both methods, power is transferred over four conductors, between the two pairs. 802.3af- and 802.3at-compliant PDs are able to accept power from either set of pairs.

Brocade PoE devices use the Endspan method, compliant with the 802.3af and 802.3at standards.

The Endspan and Midspan methods are described in more detail in the following sections.

NOTE
All 802.3af- and 802.3at-compliant power consuming devices are required to support both application methods defined in the 802.3af and 802.3at specification.

PoE endspan method

The PoE Endspan method uses the Ethernet switch ports on power sourcing equipment, such as a Brocade FastIron PoE switch, which has embedded PoE technology to deliver power over the network.

With the Endspan solution, there are two supported methods of delivering power. In Alternative A, four wires deliver data and power over the network. Specifically, power is carried over the live wire pairs that deliver data, as illustrated in Figure 8. In Alternative B, the four wires of the spare pairs are used to deliver power over the network. Brocade PoE devices support Alternative A.

The Endspan method is illustrated in Figure 8.

FIGURE 8 **PoE Endspan delivery method**

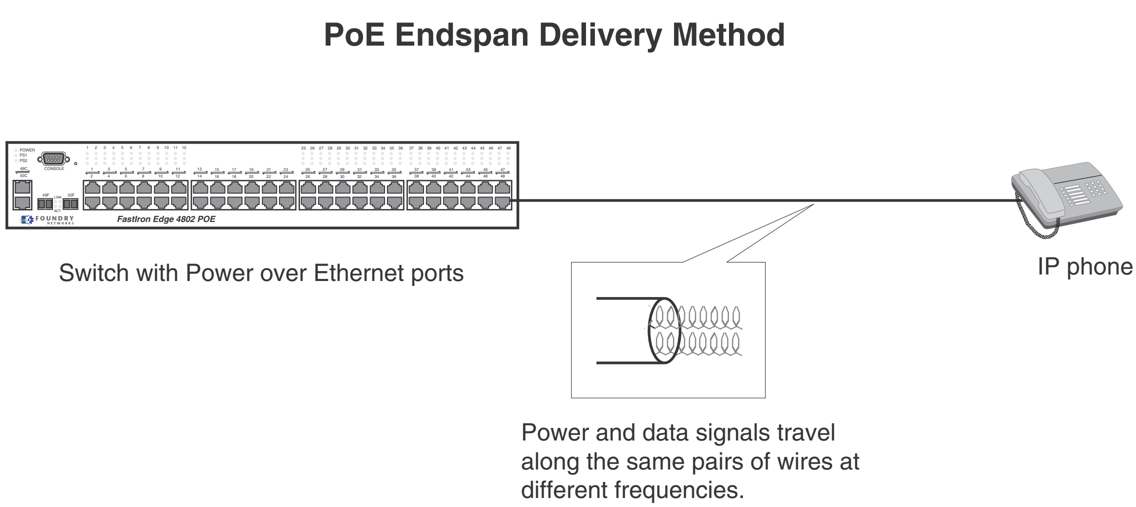

PoE midspan method

The PoE Midspan method uses an intermediate device, usually a PD, to inject power into the network. The intermediate device is positioned between the switch and the PD and delivers power over the network using the spare pairs of wires (Alternative B). The intermediate device has multiple channels (typically 6 to 24), and each of the channels has data input and a data-plus-power RJ-45 output connector.

The Midspan method is illustrated in Figure 9.

FIGURE 9 PoE Midspan delivery method

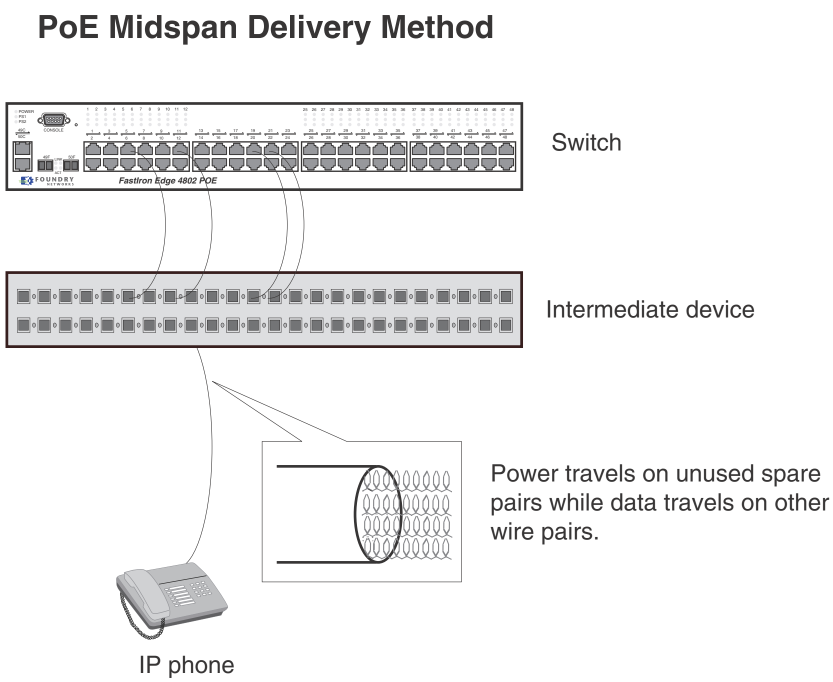

PoE autodiscovery

PoE autodiscovery is a detection mechanism that identifies whether or not an installed device is 802.3af- or 802.3at-compatible. When you plug a device into an Ethernet port that is capable of providing inline power, the autodiscovery mechanism detects whether or not the device requires power and how much power is needed. The autodiscovery mechanism also has a disconnect protection mechanism that shuts down the power once a PD has been disconnected from the network or when a faulty PD has been detected. This feature enables safe installation and prevents high-voltage damage to equipment.

PoE autodiscovery is achieved by periodically transmitting current or test voltages that can detect when a PD is attached to the network. When an 802.3af- or 802.3at-compatible device is plugged into a PoE or PoE+ port, the PD reflects test voltage back to the power sourcing device (the Brocade device), ultimately causing the power to be switched on. Devices not compatible with 802.3af do not reflect test voltage back to the power sourcing device.

Power class

A power class determines the amount of power a PD receives from a PSE. When a valid PD is detected, the Brocade PoE device performs power classification by inducing a specific voltage and measuring the current consumption of the PD. Depending on the measured current, the appropriate class is assigned to the PD. PDs that do not support classification are assigned a class of 0 (zero). Table 31 shows the different power classes and their respective power consumption needs.

TABLE 31 Power classes for PDs

Class	Usage	Power (watts) from Power Sourcing Device	
		Standard PoE	PoE+
0	default	15.4	30
1	optional	4	4
2	optional	7	7
3	optional	15.4	15.4
4	optional	15.4	30

Power specifications

The 802.3af (PoE) standard limits power to 15.4 watts (44 to 50 volts) from the power sourcing device, in compliance with safety standards and existing wiring limitations. Though limited by the 802.3af standard, 15.4 watts of power was ample for most PDs, which consumed an average of 5 to 12 watts of power (IP phones, wireless LAN access points, and network surveillance cameras each consume an average of 3.5 to 9 watts of power). The newer 802.3at (PoE+) standard nearly doubles the power, providing 30 watts (52 or 54 volts) from the power sourcing device.

The PoE power supply provides power to the PoE circuitry block, and ultimately to PoE power-consuming devices. The number of PoE power-consuming devices that one PoE power supply can support depends on the number of watts required by each power-consuming device. Each PoE power supply can provide either 1080 or 2380 watts of power, and each PoE port supports a maximum of either 15.4 or 30 watts of power per power-consuming device. For example, if each PoE power-consuming device attached to a FastIron PoE device consumes 10 watts of power, one 1080 watt power supply will power up to 108 PoE ports. You can install a second PoE power supply for additional PoE power. Power supply specifications are covered in the *Brocade FastIron X Series Chassis Hardware Installation Guide* and in the *Brocade FastIron CX Hardware Installation Guide*.

Dynamic upgrade of PoE power supplies

NOTE
This section applies to the SX 800 and SX 1600 chassis with PoE power supplies.

PoE+ requires higher power levels than standard PoE. In a chassis running software release 07.2.00 or higher, POE power supplies (SX-ACPWR-POE) are upgraded dynamically to 52 or 54 volts, depending on the maximum operating voltage the power supplies are capable of. The preferred voltage mode for PoE+ is 54 volts.

For safety reasons, all PoE power supplies installed in the chassis must operate at the same voltage mode, either 52 volts or 54 volts. The system will select the voltage mode of the power supply with the lowest supported voltage as the voltage mode for all PoE power supplies installed in the chassis. For example, in a FSX 800 chassis with one 52-volt capable PoE power supply and one 54-volt capable PoE power supply, both power supplies will be configured dynamically to operate at 52 volts.

PoE+ voltage selection occurs during each of the following events:

- When the device is powered ON or is rebooted
- When a PoE power supply is installed in the chassis

- When a PoE power supply is removed from the chassis

These events are described in detail in the following sections.

NOTE
A PoE power supply upgrade does not persist beyond a single power cycle. Therefore, an upgrade will occur automatically each time a power supply is re-inserted in the chassis.

You can use the **show inline power detail** command to display detailed information about the PoE power supplies installed in a FastIron PoE device. For more information refer to “Displaying detailed information about PoE power supplies” on page 310.

CAUTION

The SX-POE-AC-PWR power supply is designed exclusively for use with the Brocade FSX PoE devices. The power supply produces extensive power to support 802.3af and 802.3at applications. Installing the power supply in a device other than the Brocade FSX PoE device will cause extensive damage to your equipment.

Voltage selection during bootup

During bootup, the system will select the voltage mode (either 52 volts or 54 volts) of the power supply with the lowest supported voltage as the voltage mode for all PoE power supplies installed in the chassis. For example, if there is at least one power supply that supports 52 volts maximum, then all power supplies will be configured to operate at 52 volts, even if other supplies are 54 volts-capable. Once the operating voltage is applied, the system will display and log a warning message similar to the following:

```
Brocade(config)#
Power supply 1 (from left when facing front side) detected.
Power supply 1 (from left when facing front side) is up.
WARNING: PoE power supplies in slots 1 are down rev. PoE/PoE+ function will work,
 but output power may be less than 50V under worst case load.
```

If all power supplies are 54 volts-capable, then all power supplies will be configured to operate at 54 volts. In this case, the system will not display or log a warning message.

Voltage selection when a PoE power supply is installed

When a PoE power supply is hotswapped into the chassis, the system will automatically adjust the voltage to match that of the PoE power supply or supplies that are currently installed in the chassis.

The following examples describe how the voltage is selected when a PoE power supply is installed:

- If a 54 volt-capable power supply is installed in a chassis that is operating with 52 volt-capable power supplies, the newly installed power supply will be set to operate at 52 volts.
- If a 54 volt-capable power supply is installed in a chassis that is operating with 54 volt-capable power supplies, the newly installed power supply will be set to operate at 54 volts.
- If a 52 volt-capable power supply is installed in a chassis that is operating with 54 volt-capable power supplies that are actively providing power, the system will reject the newly installed power supply since it cannot safely operate with the 54 volt-capable power supplies. In this case, the 52-volt power supply will be powered OFF and an error message similar to the following will display on the console.

```
Brocade(config)#
Power supply 1 (from left when facing front side) detected.
Power supply 1 (from left when facing front side) is up.
Shutting down power supply in slot 1 because it is not compatible with the
existing PoE power supplies. Please remove and replace.
```

When the system is next reloaded, the power supply voltage will be selected as described in the section "Voltage selection during bootup" on page 296.

- If a 52 volt-capable power supply is installed in a chassis that is operating with 54 volt-capable power supplies that are *not* actively providing power, the system will configure the power supplies to operate at 52 volts. In this case, the newly installed 52-volt power supply will *not* be powered OFF and a message similar to the following will display on the console.

```
NOTE: Automatically downgraded all PoE power supplies to 52V.
```

Voltage selection when a PoE power supply is removed

If a 52 volt PoE power supply is removed from the chassis, the system will survey the remaining power supplies to determine if they are 54 volts-capable. If the remaining supplies are 54 volts-capable and the system is not currently providing power to any PDs, then the software will upgrade the voltage of all supplies to 54 volts. The system will display and log a message similar to the following:

```
NOTE: Automatically upgraded all PoE power supplies to 54V.
```

However, if the system is currently providing power to one or more PDs, the system will not upgrade the voltage level. When the system is next reloaded, the power supply voltage will be selected as described in the section "Voltage selection during bootup" on page 296.

Power over Ethernet cabling requirements

The 802.3af and 802.3at standards currently support PoE and PoE+ on 10/100/1000-Mbps Ethernet ports operating over standard Category 5 unshielded twisted pair (UTP) cable or better. If your network uses cabling categories less than Category 5, you cannot implement PoE without first upgrading your cables to Category 5 UTP cable or better.

Supported powered devices

Brocade PoE devices support a wide range of IP powered devices including the following:

- Voice over IP (VoIP) phones
- Wireless LAN access points
- IP surveillance cameras

The following sections briefly describe these IP powered devices.

VoIP

Voice over IP (VoIP) is the convergence of traditional telephony networks with data networks, utilizing the existing data network infrastructure as the transport system for both services. Traditionally, voice is transported on a network that uses circuit-switching technology, whereas data networks are built on packet-switching technology. To achieve this convergence, technology has

been developed to take a voice signal, which originates as an analog signal, and transport it within a digital medium. This is done by devices, such as VoIP telephones, which receive the originating tones and place them in UDP packets, the size and frequency of which is dependant on the coding / decoding (CODEC) technology that has been implemented in the VoIP telephone or device. The VoIP control packets use the TCP/IP format.

IP surveillance cameras

IP surveillance technology provides digital streaming of video over Ethernet, providing real-time, remote access to video feeds from cameras.

The main benefit of using IP surveillance cameras on the network is that you can view surveillance images from any computer on the network. If you have access to the Internet, you can securely connect from anywhere in the world to view a chosen facility or even a single camera from your surveillance system. By using a Virtual Private Network (VPN) or the company intranet, you can manage password-protected access to images from the surveillance system. Similar to secure payment over the Internet, images and information are kept secure and can be viewed only by approved personnel.

Installing PoE firmware

PoE firmware is stored in the PoE controller of the FastIron switch. You can install PoE firmware from the TFTP server on a FastIron switch with the CLI command. To do so, you should have a valid firmware image on the TFTP server.

NOTE
You can install PoE firmware only on one switch at a time. Therefore, to install PoE firmware on a stacking unit, you need to install it individually on every switch of the stack.

NOTE
The CLI syntax to install PoE firmware is different on FSX and FCX platforms.

FSX platform

To install PoE firmware on a FSX platform, enter a command such as the following.

```
Brocade#inline power install-firmware module 1  tftp 10.120.54.161
fsx_poe_07400.fw
```

Syntax: **inline power install-firmware module** *slot* **tftp** *ip-address filename*

Slot refers to the slot of the PoE module.

ip-address refers to the IP address of the TFTP server.

Filename refers to the name of the file, including the pathname.

FCX and ICX platforms

To install PoE firmware on FCX and ICX platforms, enter a command such as the following.

```
Brocade#inline power install-firmware stack-unit 1  tftp 10.120.54.161
fcx_poeplus_07400.fw
```

Syntax: **inline power install-firmware stack-unit** *unit-number* **tftp** *ip-address filename*

Stack-unit refers to the unit-id of the switch. If the switch is not a part of the stack, the unit number will be the default value. The default value for stack-unit is 1.

ip-address refers to the IP address of the tftp server.

Filename refers to the name of the file, including the pathname.

If you want to install firmware on a stack, you need to install firmware on one switch at a time with the above command.

Image file types

This section lists the PoE firmware file types supported and the procedure to install them on the FCX, ICX, and FSX devices.

TABLE 32 PoE Firmware files

Product	PoE Firmware[1]
FSX 800 FSX 1600	fsx_poe_07400.fw
FSX 800 with SX-FI648PP module FSX 1600 with SX-FI648PP module	fsx_poeplus_07400.fw
FCX ICX 6610	fcx_poeplus_07400.fw
ICX 6430 ICX 6450	icx64XX_poeplus_07400.fw

1. The firmware files are specific to these devices only and are not interchangeable. For example, you cannot load FCX PoE firmware on a FSX device, and vice versa.

Installing PoE firmware

1. Place the PoE firmware on a TFTP server to which the Brocade device has access.
2. Copy the PoE firmware from the TFTP server into the switch. To do so, enter a command such as the following.

```
Brocade#inline power install-firmware module 1  tftp 10.120.54.161
fsx_poe_07400.fw
```

The process of PoE installation begins. You should see output similar to the following.

```
PoE Info: Loading firmware from TFTP file fsx_poe_07400.fw........

PoE Info: Firmware in PoE module(s) in slot 1 will be installed now.
PoE Warning: Upgrading firmware in slot 1....DO NOT HOTSWAP OR POWER DOWN THE
MODULE.

PoE Info: FW Download on slot 1 module 1...(re)sending download command...
PoE Info: FW Download on slot 1 module 1...TPE response received.

PoE Info: FW Download on slot 1 module 1...(re)sending erase command...
PoE Info: FW Download on slot 1 module 1...erase command...accepted.
PoE Info: FW Download on slot 1 module 1...erasing firmware memory...
PoE Info: FW Download on slot 1 module 1...erasing firmware memory...completed

PoE Info: FW Download on slot 1 module 1...(re)sending program command...
PoE Info: FW Download on slot 1 module 1...sending program command...accepted.
```

```
PoE Info: FW Download on slot 1 module 1...programming firmware...takes ~ 10
minutes....

PoE Info: Firmware Download on slot 1.....10 percent completed.
PoE Info: Firmware Download on slot 1.....20 percent completed.
PoE Info: Firmware Download on slot 1.....30 percent completed.
PoE Info: Firmware Download on slot 1.....40 percent completed.
PoE Info: Firmware Download on slot 1.....50 percent completed.
PoE Info: Firmware Download on slot 1.....60 percent completed.
PoE Info: Firmware Download on slot 1.....70 percent completed.
PoE Info: Firmware Download on slot 1.....80 percent completed.
PoE Info: Firmware Download on slot 1.....90 percent completed.
PoE Info: Firmware Download on slot 1.....100 percent completed.
PoE Info: FW Download on slot 1 module 1...programming firmware...completed.
```

3. After downloading the firmware into the controller, the controller resets and reboot with the new PoE firmware, You should see output similar to the following.

```
PoE Info: FW Download on slot 1 module 1...upgrading firmware...completed.
Module will be reset.
PoE Info: Resetting module 1 in slot 1....
PoE Info: Resetting module 1 in slot 1....completed.

PoE Info: Programming Brocade defaults on module 1 in slot 1.....
PoE Info: Programming Brocade defaults. Step 1: Writing port defaults on
module 1 in slot 1....
PoE Info: Programming Brocade Defaults: Step 2: Writing PM defaults on module
1 in slot 1.
PoE Info: Programming Brocade defaults. Step 3: Writing user byte 0xf0 on
module 1 in slot 1.
PoE Info: Programming Brocade defaults. Step 4: Saving settings on module 1 in
slot 1.
PoE Info: Programming Brocade defaults on module 1 in slot 1.....completed.
```

NOTE
If you are attempting to transfer a file using TFTP but have received an error message, refer to "Diagnostic error codes and remedies for TFTP transfers" on page 94.

PoE and CPU utilization

Depending on the number of PoE-configured ports that have powered power devices, there may be a slight and noticeable increase of up to 15 percent in CPU utilization. In typical scenarios, this is normal behavior for PoE and does not affect the functionality of other features on the switch.

Enabling and disabling Power over Ethernet

To enable a port to receive inline power for power consuming devices, enter commands such as the following.

```
Brocade#configure terminal
Brocade(config)# interface ethernet 1/1
Brocade(config-if-e1000-1/1)# inline power
```

After entering the above commands, the console displays the following message.

```
Brocade(config-if-e1000-1/1)#PoE Info: Power enabled on port 1/1.
```

Syntax: [**no**] **inline power**

Use the **no** form of the command to disable the port from receiving inline power.

NOTE
Inline power should not be configured between two switches as it may cause unexpected behavior.

NOTE
FastIron PoE and PoE+ devices can automatically detect whether or not a power consuming device is 802.3af- or 802.3at-compliant.

Disabling support for PoE legacy power-consuming devices

Brocade PoE devices automatically support most legacy power consuming devices (devices not compliant with 802.3af 802.3at), as well as all 802.3af- and 802.3at-compliant devices. If desired, you can disable and re-enable support for legacy PoE power consuming devices on a global basis (on the entire device) or on individual slots (chassis devices only). When you disable legacy support, 802.3af- and 802.3at-compliant devices are not affected.

To disable support for legacy power consuming devices on a non-stackable device, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)# no legacy-inline-power
```

To disable support for legacy power consuming devices on a stackable device, enter the following command at the stack unit CONFIG level of the CLI.

```
Brocade(config-unit-2)# no legacy-inline-power
```

On chassis devices, you can disable support for legacy power consuming devices per slot. To disable legacy support on all ports in slot 2, enter the following command at the global CONFIG level of the CLI.

```
Brocade(config)# no legacy-inline-power 2
```

NOTE
The **no legacy-inline-power** command does not require a software reload if it is entered prior to connecting the PDs. If the command is entered after the PDs are connected, the configuration must be saved (**write memory**) and the software reloaded after the change is placed into effect.

Syntax: [**no**] **legacy-inline-power** [*slotnum*]

NOTE
By default, the inline-power command reserves 30 watts.

To re-enable support for legacy power consuming devices after it has been disabled, enter the **legacy-inline-power** command (without the **no** parameter).

The *slotnum* variable is required on chassis devices when disabling or re-enabling legacy support on a slot.

Use the **show run** command to view whether support for PoE legacy power consuming devices is enabled or disabled.

Enabling the detection of PoE power requirements advertised through CDP

Many power consuming devices, such as Cisco VoIP phones and other vendors' devices, use the Cisco Discovery Protocol (CDP) to advertise their power requirements to power sourcing devices, such as Brocade PoE devices. Brocade power sourcing equipment is compatible with Cisco and other vendors' power consuming devices; they can detect and process power requirements for these devices automatically.

NOTE
If you configure a port with a maximum power level or a power class for a power consuming device, the power level or power class will take precedence over the CDP power requirement. Therefore, if you want the device to adhere to the CDP power requirement, do not configure a power level or power class on the port.

Command syntax for PoE power requirements

To enable the Brocade device to detect CDP power requirements, enter the following commands.

```
Brocade# configure terminal
Brocade(config)# cdp run
```

Syntax: [**no**] **cdp run**

Use the **no** form of the command to disable the detection of CDP power requirements.

Setting the maximum power level for a PoE power-consuming device

When PoE is enabled on a port to which a power consuming device or PD is attached, by default, a Brocade PoE device will supply 15.4 watts of power at the RJ-45 jack, minus any power loss through the cables. A PoE+ device will supply either 15.4 or 30 watts of power (depending on the type of PD connected to the port), minus any power loss through the cables. For example, a PoE port with a default maximum power level of 15.4 watts will receive a maximum of 12.95 watts of power after 2.45 watts of power loss through the cable. This is compliant with the IEEE 802.3af and 802.3at specifications for delivering inline power. Devices that are configured to receive less PoE power, for example, 4.0 watts of power, will experience a lower rate of power loss through the cable.

If desired, you can manually configure the maximum amount of power that the Brocade PoE device will supply at the RJ-45 jack.

Setting power levels configuration note

Consider the following when enabling this feature:

- There are two ways to configure the power level for a PoE or PoE+ power consuming device. The first method is discussed in this section. The other method is provided in the section "Setting the power class for a PoE power- consuming device" on page 304. For each PoE port, you can configure either a maximum power level or a power class. You cannot configure both. You can, however, configure a maximum power level on one port and a power class on another port.
- The Brocade PoE or PoE+ device will adjust the power on a port only if there are available power resources. If power resources are not available, the following message will display on the console and in the Syslog:

```
PoE: Failed power allocation of 30000 mwatts on port 1/1/21. Will retry when
more power budget.
```

Configuring power levels command syntax

To configure the maximum power level for a power consuming device, enter commands such as the following.

```
Brocade#configure terminal
Brocade(config)# interface ethernet 1/1
Brocade(config-if-e1000-1/1)# inline power power-limit 14000
```

These commands enable inline power on interface ethernet 1 in slot 1 and set the PoE power level to 14,000 milliwatts (14 watts).

Syntax: **inline power power-limit** *power level*

The *power level* variable is the maximmum power level in number of milliwatts. The following values are supported:

- PoE – Enter a value from 1000 through 15,400. The default is 15,400.
- PoE+ – Enter a value from 1000 through 30,000. The default is 30,000.

NOTE
Do not configure a power level higher than 15,400 for standard PoE PDs, which support a maximum of 15,400 milliwatts. Setting the power level higher than 15,400 could damage the PD.

For information about resetting the maximum power level, refer to "Resetting PoE parameters" on page 307.

Setting the power class for a PoE power-consuming device

A power class specifies the maximum amount of power that a Brocade PoE or PoE+ device will supply to a power consuming device. Table 33 shows the different power classes and their respective maximum power allocations.

TABLE 33 Power classes for PDs

Class	Usage	Power (watts) from Power Sourcing Device	
		Standard PoE	PoE+
0	default	15.4	30
1	optional	4	4
2	optional	7	7
3	optional	15.4	15.4
4	optional	15.4	30

Setting the power class configuration notes

Consider the following points when setting the power class for a PoE power-consuming device.

- The power class sets the maximum power level for a power consuming device. Alternatively, you can set the maximum power level as instructed in the section "Setting the maximum power level for a PoE power- consuming device" on page 302. For each PoE port, you can configure either a power class or a maximum power level. You cannot configure both. You can, however, configure a power level on one port and a power class on another port.
- The power class includes any power loss through the cables. For example, a PoE port with a power class of 3 (15.4 watts) will receive a maximum of 12.95 watts of power after 2.45 watts of power loss through the cable. This is compliant with the IEEE 802.3af and 802.3at specifications for delivering inline power. Devices that are configured to receive less PoE power, for example, class 1 devices (4.0 watts), will experience a lower rate of power loss through the cable.
- The Brocade PoE or PoE+ device will adjust the power on a port only if there are available power resources. If power resources are not available, the following message will display on the console and in the Syslog:

  ```
  PoE: Failed power allocation of 30000 mwatts on port 1/1/21. Will retry when
  more power budget.
  ```

Setting the power class command syntax

To configure the power class for a PoE power consuming device, enter commands such as the following.

```
Brocade# configure terminal
Brocade(config)# interface ethernet 1/1
Brocade(config-if-e1000-1/1)# inline power power-by-class 2
```

These commands enable inline power on interface ethernet 1 in slot 1 and set the power class to 2.

Syntax: **inline power power-by-class** *class value*

The *class value* variable is the power class. Enter a value between 0 and 4. The default is 0. Table 33 shows the different power classes and their respective maximum power allocations.

NOTE
Do not configure a class value of 4 on a PoE+ port on which a standard PoE PD is connected. Standard PoE PDs support a maximum of 15.4 watts. Setting the power class value to 4 (30 watts) could damage the PD.

For information about resetting the power class, refer to “Resetting PoE parameters” on page 307.

Setting the power budget for a PoE interface module

By default, each PoE and PoE+ interface module has a maximum power budget of 65535 watts. If desired, you can change the amount of power allocated to each PoE and PoE+ interface module installed in the chassis. To do so, enter a command such as the following.

```
Brocade(config)# inline power budget 150000 module 7
```

This command allocates 150000 milliwatts (150 watts) to the PoE interface module in slot 7. The command takes effect immediately. The results are displayed in the “power budget” column in the **show inline power detail** output. The configuration (`inline power budget 150000 module 7`) is displayed in the **show running-config** output.

Syntax: **inline power budget** *num* **module** *slot*

The *num* variable is the number of milliwatts to allocate to the module. Enter a value from 0 through 65535000.

The *slot* variable specifies the where the PoE or PoE+ module resides in the chassis.

Setting the inline power priority for a PoE port

Each PoE power supply can provide either 1080 or 2380 watts of power, and each PoE port receives a maximum of 15.4 watts of power per PoE power-consuming device, or a maximum of 30 watts of power per PoE+ power-consuming device, minus any power loss through the cable. The power capacity of one or two PoE power supplies is shared among all PoE power consuming devices attached to the FastIron PoE device.

In a configuration where PoE power consuming devices collectively have a greater demand for power than the PoE power supply or supplies can provide, the FastIron PoE device must place the PoE ports that it cannot power in *standby* or *denied* mode (waiting for power) until the available power increases. The available power increases when one or more PoE ports are powered down, or, if applicable, when an additional PoE power supply is installed in the FastIron PoE device.

When PoE ports are in *standby* or *denied* mode (waiting for power) and the FastIron PoE device receives additional power resources, by default, the device will allocate newly available power to the standby ports in priority order, with the highest priority ports first, followed by the next highest priority ports, and so on. Within a given priority, standy ports are considered in ascending order, by slot number then by port number, provided enough power is available for the ports. For example,

PoE port 1/11 should receive power before PoE port 2/1. However, if PoE port 1/11 needs 12 watts of power and PoE port 2/1 needs 10 watts of power, and 11 watts of power become available on the device, the FastIron PoE device will allocate the power to port 2/1 because it does not have sufficient power for port 1/11.

You can configure an *inline power priority* on PoE ports, whereby ports with a higher inline power priority will take precedence over ports with a low inline power priority. For example, if a new PoE port comes online and the port is configured with a high priority, if necessary (if power is already fully allocated to power consuming devices), the FastIron PoE device will remove power from a PoE port or ports that have a lower priority and allocate the power to the PoE port that has the higher value.

Ports that are configured with the same inline power priority are given precedence based on the slot number and port number in ascending order, provided enough power is available for the port. For example, if both PoE port 1/2 and PoE port 2/1 have a high inline power priority value, PoE port 1/2 will receive power before PoE port 2/1. However, if PoE port 1/2 needs 12 watts of power and PoE port 2/1 needs 10 watts of power, and 11 watts of power become available on the device, the FastIron PoE device will allocate the power to PoE port 2/1 because it does not have sufficient power for port 1/2. By default, all ports are configured with a low inline power priority.

Command syntax for setting the inline power priority for a PoE port

To configure an inline power priority for a PoE port on a FastIron PoE device, enter commands such as the following.

```
Brocade#configure terminal
Brocade(config)# interface ethernet 1/1
Brocade(config-if-e1000-1/1)# inline power priority 2
```

These commands enable inline power on interface ethernet 1 in slot 1 and set the inline power priority level to high.

Syntax: [**no**] **inline power priority** *priority num*

The **priority** *priority num* parameter is the inline power priority number. The default is 3 (low priority). You can specify one of the following values:

- 3 – Low priority
- 2 – High priority
- 1 – Critical priority

Use the **inline power** command (without a priority number) to reset a port priority to the default (low) priority.

Use the **no inline power** command to disable the port from receiving inline power.

For information about resetting the inline power priority, refer to “Resetting PoE parameters” on page 307.

To view the inline power priority for all PoE ports, issue the **show inline power** command at the Privileged EXEC level of the CLI. Refer to “Displaying PoE operational status” on page 307.

Resetting PoE parameters

NOTE
Resetting PoE parameters applies to the FastIron X Series PoE chassis.

You can override or reset PoE port parameters including power priority, power class, and maximum power level. To do so, you must specify each PoE parameter in the CLI command line. This section provides some CLI examples.

Example 1—Changing a PoE port power priority from high to low

To change a PoE port power priority from high to low (the default value) and keep the current maximum configured power level of 3000, enter commands such as the following.

```
Brocade# configure terminal
Brocade(config)# interface ethernet 1/1
Brocade(config-if-e1000-1/1)# inline power priority 3 power-limit 3000
```

You must specify both the inline power priority and the maximum power level (**power-limit** command), even though you are keeping the current configured maximum power level at 3000. If you do not specify the maximum power level, the device will apply the default value. Also, you must specify the inline power priority before specifying the power limit.

Example 2—Changing a port power class from 2 to 3

To change a port power class from 2 (7 watts maximum) to 3 (15.4 watts maximum) and keep the current configured power priority of 2, enter commands such as the following.

```
Brocade#configure terminal
Brocade(config)# interface ethernet 1/1
Brocade(config-if-e1000-1/1)# inline power priority 2 power-by-class 3
```

You must specify both the power class and the inline power priority, even though you are not changing the power priority. If you do not specify the power priority, the device will apply the default value of 3 (low priority). Also, you must specify the inline power priority before specifying the power class.

Displaying Power over Ethernet information

This section lists the CLI commands for viewing PoE information.

Displaying PoE operational status

The **show inline power** command displays operational information about Power over Ethernet.

You can view the PoE operational status for the entire device, for a specific PoE module only, or for a specific interface only. In addition, you can use the **show inline power detail** command to display in-depth information about PoE power supplies.

The following shows an example of the **show inline power** display output on a Brocade PoE device.

```
Brocade#show inline power

Power Capacity:         Total is 2160000 mWatts. Current Free is 18800 mWatts.

Power Allocations:      Requests Honored 769 times

 ... some lines omitted for brevity...

 Port   Admin   Oper    ---Power(mWatts)---  PD Type  PD Class  Pri  Fault/
        State   State   Consumed  Allocated                          Error
--------------------------------------------------------------------------
 4/1    On      On          5070       9500  802.3af  n/a         3  n/a
 4/2    On      On          1784       9500  Legacy   n/a         3  n/a
 4/3    On      On          2347       9500  802.3af  n/a         3  n/a
 4/4    On      On          2441       9500  Legacy   n/a         3  n/a
 4/5    On      On          6667       9500  802.3af  Class 3     3  n/a
 4/6    On      On          2723       9500  802.3af  Class 2     3  n/a
 4/7    On      On          2347       9500  802.3af  n/a         3  n/a
 4/8    On      On          2347       9500  802.3af  n/a         3  n/a
 4/9    On      On          2347       9500  802.3af  n/a         3  n/a
 4/10   On      On          4976       9500  802.3af  Class 3     3  n/a
 4/11   On      On          4882       9500  802.3af  Class 3     3  n/a
 4/12   On      On          4413       9500  802.3af  Class 1     3  n/a
 4/13   On      On          7793       9500  802.3af  n/a         3  n/a
 4/14   On      On          7512       9500  802.3af  n/a         3  n/a
 4/15   On      On          8075       9500  802.3af  n/a         3  n/a
 4/16   On      On          4131       9500  802.3af  Class 1     3  n/a
 4/17   On      On          2347       9500  802.3af  n/a         3  n/a
 4/18   On      Off            0       9500  n/a      n/a         3  n/a
 4/19   On      On          5352       9500  Legacy   n/a         3  n/a
 4/20   On      On          7981       9500  802.3af  n/a         3  n/a
 4/21   On      On         12958      13000  802.3af  Class 3     3  n/a
 4/22   On      On         12958      13000  802.3af  Class 3     3  n/a
 4/23   On      On         13052      13000  802.3af  Class 3     3  n/a
 4/24   On      On         12864      13000  802.3af  Class 3     3  n/a
--------------------------------------------------------------------------
 Total                    137367     242000

 ... some lines omitted for brevity...

 Grand Total             1846673    2127400
```

Syntax: show inline power [*port*]

Table 34 provides definitions for the **show inline power** command.

TABLE 34 Field definitions for the show inline power command

Column	Definition
Power Capacity	The total PoE power supply capacity and the amount of available power (current free) for PoE power consuming devices. Both values are shown in milliwatts.
Power Allocations	The number of times the FSX fulfilled PoE requests for power.
Port	The slot number and port number.
Admin State	Specifies whether or not Power over Ethernet has been enabled on the port. This value can be one of the following: • On – The **inline power** command was issued on the port. • Off – The **inline power** command has not been issued on the port.
Oper State	Shows the status of inline power on the port. This value can be one of the following: • On – The PoE power supply is delivering inline power to the PD. • Off – The PoE power supply is not delivering inline power to the PD. • Denied – The port is in standby mode (waiting for power) because the FSX does not currently have enough available power for the port. NOTE: When you enable a port using the CLI, it may take 12 or more seconds before the operational state of that port is displayed correctly in the **show inline power** output.
Power Consumed	The number of current, actual milliwatts that the PD is consuming.
Power Allocated	The number of milliwatts allocated to the port. This value is either the default or configured maximum power level, or the power class that was automatically detected by the device.

PD Type	The type of PD connected to the port. This value can be one of the following:
	• 802.3at – The PD connected to this port is 802.3at-compliant.802.3af – The PD connected to this port is 802.3af-compliant.
	• Legacy – The PD connected to this port is a legacy product (not 802.3af-compliant).
	• N/A – Power over Ethernet is configured on this port, and one of the following is true:
	  • The device connected to this port is a non-powered device.
	  • No device is connected to this port.
	  • The port is in *standby* or *denied* mode (waiting for power).

PD Class	Determines the maximum amount of power a PD receives. Table 33 shows the different power classes and their respective maximum power allocations. This field can also be “Unknown”, meaning the device attached to the port cannot advertise its power class. NOTE: If an 802.3at PD with a class 4 value is connected, the Brocade FastIron switch will not be able to do the power negotiation since these switches cannot handle the 802.3at LLDP.
Pri	The port *in-line power priority*, which determines the order in which the port will receive power while in *standby* mode (waiting for power). Ports with a higher priority will receive power before ports with a low priority. This value can be one of the following: • 3 – Low priority • 2 – High priority • 1 – Critical priority

TABLE 34 Field definitions for the show inline power command (Continued)

Column	Definition
Fault/Error	If applicable, this is the fault or error that occurred on the port. This value can be one of the following: • critical temperature – The PoE chip temperature limit rose above the safe operating level, thereby powering down the port. • detection failed - discharged capacitor – The port failed capacitor detection (legacy PD detection) because of a discharged capacitor. This can occur when connecting a non-PD on the port. • detection failed - out of range capacitor – The port failed capacitor detection (legacy PD detection) because of an out-of-range capacitor value. This can occur when connecting a non-PD on the port. • internal h/w fault – A hardware problem has hindered port operation. • lack of power – The port has shut down due to lack of power. • main supply voltage high – The voltage was higher than the maximum voltage limit, thereby tripping the port. • main supply voltage low – The voltage was lower than the minimum voltage limit, thereby tripping the port. • overload state – The PD consumes more power than the maximum limit configured on the port, based on the default configuration, user configuration, or CDP configuration. • over temperature – The port temperature rose above the temperature limit, thereby powering down the port. • PD DC fault – A succession of underload and overload states, or a PD DC/DC fault, caused the port to shutdown. • short circuit – A short circuit was detected on the port delivering power. • underload state – The PD consumes less power than the minimum limit specified in the 802.3af standard. • voltage applied from ext src – The port failed capacitor detection (legacy PD detection) because the voltage applied to the port was from an external source.
Total	The total power in milliwatts being *consumed* by all PDs connected to the Interface module, and the total power in milliwatts *allocated* to all PDs connected to the Interface module.
Grand Total	The total number of current, actual milliwatts being *consumed* by all PDs connected to the FastIron PoE device, and the total number of milliwatts *allocated* to all PDs connected to the FastIron PoE device.

Displaying detailed information about PoE power supplies

The **show inline power detail** command displays detailed operational information about the PoE power supplies in Brocade PoE switches. The command output differs on FCX POE+ switches compared to FastIron X Series switches.

To following is an example of the **show inline power detail** command output on an FCX POE+ switch.

```
Brocade#FCX#show inline power detail

Power Supply Data On stack 1:
++++++++++++++++++
Power Supply #1:
        Max Curr:       7.5 Amps
        Voltage:        54.0 Volts
        Capacity:       410 Watts

POE Details Info. On Stack 1 :

General PoE Data:
+++++++++++++++++

Firmware
Version
--------
02.1.0

Cumulative Port State Data:
+++++++++++++++++++++++++++

#Ports    #Ports     #Ports   #Ports    #Ports        #Ports     #Ports
Admin-On  Admin-Off  Oper-On  Oper-Off  Off-Denied    Off-No-PD  Off-Fault
-------------------------------------------------------------------------
45        3          0        48        0             45         0

Cumulative Port Power Data:
+++++++++++++++++++++++++++

#Ports  #Ports  #Ports        Power       Power
Pri: 1  Pri: 2  Pri: 3  Consumption  Allocation
-----------------------------------------------
0       0       45          0.0   W     0.0   W

Power Supply Data On stack 2:
++++++++++++++++++

Power Supply Data:
++++++++++++++++++

Power Supply #1:
        Max Curr:       7.5 Amps
        Voltage:        54.0 Volts
        Capacity:       410 Watts

POE Details Info. On Stack 2 :

General PoE Data:
+++++++++++++++++
Firmware
Version
--------
02.1.0
```

```
... continued from previous page...

Cumulative Port State Data:
+++++++++++++++++++++++++++

#Ports    #Ports     #Ports   #Ports    #Ports       #Ports     #Ports
Admin-On  Admin-Off  Oper-On  Oper-Off  Off-Denied   Off-No-PD  Off-Fault
-------------------------------------------------------------------------
20        4          0        24        0            20         0

Cumulative Port Power Data:
+++++++++++++++++++++++++++
#Ports  #Ports  #Ports        Power       Power
Pri: 1  Pri: 2  Pri: 3  Consumption  Allocation
-----------------------------------------------
20      0       0           0.0   W     0.0   W

Power Supply Data On stack 3:
++++++++++++++++++

Power Supply #1:
        Max Curr:       7.5 Amps
        Voltage:        54.0 Volts
        Capacity:       410 Watts

POE Details Info. On Stack 3 :

General PoE Data:
+++++++++++++++++

Firmware
Version
--------
02.1.0

Cumulative Port State Data:
+++++++++++++++++++++++++++

#Ports    #Ports     #Ports   #Ports    #Ports       #Ports     #Ports
Admin-On  Admin-Off  Oper-On  Oper-Off  Off-Denied   Off-No-PD  Off-Fault
-------------------------------------------------------------------------
22        2          0        24        0            22         0

Cumulative Port Power Data:
+++++++++++++++++++++++++++

#Ports  #Ports  #Ports        Power       Power
Pri: 1  Pri: 2  Pri: 3  Consumption  Allocation
-----------------------------------------------
0       10      12          0.0   W     0.0   W
```

To following is an example of the **show inline power detail** command output on a FastIron X Series PoE switch.

```
Brocade#show inline power detail

Power Supply Data:
++++++++++++++++++

PoE+ Max Operating Voltage: 54 V

Power Supply #1:
        Model Number:   32004000
        Serial Number:  093786124716
        Firmware Ver:   1.6
        Test Date:      9/12/09 (mm/dd/yy)
        H/W Status:     807
        Max Curr:       50.0 Amps
        Voltage:        54.0 Volts
        Capacity:       2500 Watts
        PoE Capacity:   2260 Watts
        Consumption:    2095 Watts


General PoE Data:
+++++++++++++++++

Slot  Firmware
      Version
--------------
3     Device 1: 02.1.0  Device 2: 02.1.0
4     Device 1: 02.1.0  Device 2: 02.1.0
6     02.1.0
7     Device 1: 02.1.0  Device 2: 02.1.0
8     02.1.0


Cumulative Port State Data:
+++++++++++++++++++++++++++

Slot  #Ports    #Ports     #Ports   #Ports    #Ports       #Ports     #Ports
      Admin-On  Admin-Off  Oper-On  Oper-Off  Off-Denied   Off-No-PD  Off-Fault
--------------------------------------------------------------------------------
3     48        0          48       0         0            0          0
4     48        0          48       0         0            0          0
6     24        0          0        24        0            24         0
7     48        0          4        44        44           0          0
8     24        0          0        24        0            24         0
--------------------------------------------------------------------------------
Total:192       0          100      92        44           48         0
```

... continued on next page...

```
... continued from previous page...

Cumulative Port Power Data:
+++++++++++++++++++++++++++

Slot  #Ports  #Ports  #Ports       Power        Power         Power
      Pri: 1  Pri: 2  Pri: 3  Consumption  Allocation        Budget
-------------------------------------------------------------------
3     0       0       48        513.90  W    739.200 W    65535.0 W
4     0       0       48       1346.497 W   1440.0   W    65535.0 W
6     0       0       24          0.0   W      0.0   W    65535.0 W
7     0       0       48         43.72  W     61.600 W    65535.0 W
8     0       0       24          0.0   W      0.0   W    65535.0 W
-------------------------------------------------------------------
Total:0       0       192      1902.659 W   2240.800 W   327675.0 W
```

Syntax: show inline power detail

Table 35 provides definitions for the statistics displayed in the **show inline power detail** command.

TABLE 35 Field definitions for the show inline power detail command

Column	Definition
Power supply data	
PoE+ Max Operating Voltage	This field is applicable to FastIron PoE+ chassis devices only. It displays the maximum operating voltage supported by the PoE power supply. Possible values are: • 52 V • 54 V
Model Number	The manufacturing part number of the PoE power supply. Possible values are: • 32016000 • 32007000
Serial Number	The serial number of the PoE power supply, for example, AA100730213.
Firmware Ver	The PoE power supply firmware version.
Test Date	The PoE power supply firmware test date in the format mm/dd/yyyy.
H/W Status	The PoE power supply hardware status code. This field is used by Brocade Technical Support for troubleshooting.
Max Curr	The PoE power supply maximum current capacity.
Voltage	The PoE power supply current input voltage.
Capacity	The PoE power supply total power capacity (in watts).
PoE Capacity	The PoE power supply PoE power capacity (in watts).
Consumption	The total number of watts consumed by PoE power consuming devices and PoE modules in the system, plus any internal or cable power loss. **NOTE:** Under thelower total inline power consumption level by Powered Devices (PDs) on FastIron SX devices, the power consumption displayed by the power supply units (PSUs) is inaccurately displayed as lower than the actual power consumption of the PSUs due to the sensitivity limitations of power supply measurements.
General PoE data	
Slot	The Interface module / slot number.

TABLE 35 Field definitions for the **show inline power detail** command (Continued)

Column	Definition
Firmware Version	The Interface module / slot number firmware version.
Cumulative port state data **NOTE:** When you enable a port using the CLI, it may take 12 or more seconds before the operational state of that port is displayed correctly in the **show inline power** output.	
Slot	The Interface module / slot number.
#Ports Admin-On	The number of ports on the Interface module on which the **inline power** command was issued.
#Ports Admin-Off	The number of ports on the Interface module on which the **inline power** command was not issued.
#Ports Oper-On	The number of ports on the Interface module that are receiving inline power from the PoE power supply.
#Ports Oper-Off	The number of ports on the Interface module that are not receiving inline power from the PoE power supply.
#Ports Off-Denied	The number of ports on the Interface module that were denied power because of insufficient power.
#Ports Off-No-PD	The number of ports on the Interface module to which no PDs are connected.
#Ports Off-Fault	The number of ports on the Interface module that are not receiving power because of a subscription overload.
Total	The totals for all of the fields in the **Cumulative Port State Data** report.
Cumulative port power data	
Slot	The Interface module / slot number.
#Ports Pri: 1	The number of PoE ports on the Interface module that have a PoE port priority of 1.
#Ports Pri: 2	The number of PoE ports on the Interface module that have a PoE port priority of 2.
#Ports Pri: 3	The number of PoE ports on the Interface module that have a PoE port priority of 3.
Power Consumption	The total number of watts consumed by PoE power consuming devices, plus any cable loss.
Power Allocation	The number of watts allocated to the Interface module PoE ports. This value is the sum of the ports' default or configured maximum power levels, or power classes automatically detected by the FastIron PoE device.
Power Budget	The power budget allocated to the slot. The default value is 65535 watts. Any other value indicates that the power budget was configured using the CLI command **inline power budget**.
Total	The totals for all of the fields in the **Cumulative Port Power Data** report.

System Monitoring

Chapter 13

[Table 36](#) lists the system monitoring (sysmon) features supported on Brocade FastIron devices. These features are supported in the Layer 2 and full Layer 3 software images, except where explicitly noted.

TABLE 36 Supported system monitoring features

Feature	FSX 800 FSX1600[1]	FCX	ICX 6610	ICX 6430	ICX 6450
Fabric Adaptor errors	Yes	No	No	No	No
Packet Processor errors	Yes	No	No	No	No
Cross Bar errors	Yes	No	No	No	No
Link errors	No	Yes	Yes	Yes	Yes
ECC errors	No	Yes	Yes	Yes	Yes
Sysmon Enhancements	Yes	No	No	No	No

1. Third generation modules

This chapter contains the following sections:

Overview of system monitoring

System monitoring (sysmon) is a utility that runs as a background process and monitors connections and components of the device for specific errors and logs them. It has a default policy that controls the parameters that are monitored and actions to be taken if a fault is detected. These policies include the type of errors, the threshold for errors to be logged, and the frequency of checking for errors. You can use the CLI commands to configure these policies.

The sysmon utility monitors the hardware error registers to identify errors and failures. You can configure the sysmon timer to define how frequently the sysmon utility queries the hardware error registers. The data generated by the sysmon utility is written to either the sysmon internal log or to the syslog.

Sysmon starts the timer based on the specified timer setting, with the default value as three minutes. After the interval specified by the timer, the utility checks the hardware error registers. If the sysmon utility detects an error in a hardware error register, it increments the relevant error count by 1. Otherwise, it restarts the timer and waits for the given interval. Hardware error registers are cleared when read, so after Sysmon reads the value, they are reset to zero.

Sysmon checks the value of the error counters it maintains and the values specified in the sysmon threshold. If the value of the error counters exceeds the matching threshold, it takes the action specified (logs internally or to the syslog). Otherwise, it restarts the timer and waits for the specified interval before checking for errors again.

To ensure that logging repeating errors does not cause the logs to overflow, you can specify a back-off value that allows the utility to skip the specified number of error instances before logging again. If the error count is smaller than the specified log back-off value, the utility logs the error to the internal log or syslog, restarts the timer and waits for the specified interval before checking for errors again.

Configuration notes and feature limitations

- While system monitoring is supported on all FastIron devices, the types of errors monitored vary according to devices. On FSX devices, the sysmon utility monitors the following for errors:
  - Fabric Adapter (FA) for processing and link errors.
  - Cross Bar (XBAR) or Switch Fabric Module (SFM) for processing and link errors.
  - Packet processor (PP) for link errors.

  On FCX and ICX devices, the sysmon utility monitors the following errors:
  - Link errors.
  - ECC errors.
- By default, system monitoring starts on system boot up and runs in the background every three minutes. You can configure, disable, or enable, the time interval through the CLI; however, if you define the system monitoring interval at the global level, this value overrides the individual settings. Valid range for the sysmon timer is 1 to 60 minutes.
- You can define a system monitoring threshold that is defined as N/W, where N is the number of error events in a specified window (W) of consecutive polling periods. When the threshold is reached, the action that is defined is performed. The threshold enables the sysmon utility to ignore random errors that occur because of corrupted data coming in to the device, and perform the action only for errors generated because of device failure. A threshold of 1/W means no threshold.
- You can choose the log action as either to the internal sysmon buffer or to the syslog. If you choose the internal sysmon buffer, logs that are written beyond the limit of the sysmon buffer rolls over. On the other hand, if you choose logging to syslog, messages are sent to the configured syslog servers.

Configure system monitoring

You can use the following commands at the privileged EXEC level to globally configure the sysmon utility:

- disable system-monitoring all

- enable system-monitoring all
- sysmon timer

In addition, you can enable or disable system monitoring for each event type from the CLI, with each event type having separate threshold and log back off values.

disable system-monitoring all

Disables system monitoring at the global level for all types.

Syntax **disable system-monitoring all**

Command Default By default, system monitoring is enabled.

Command Modes Privileged EXEC mode.

Usage Guidelines Disabling sysmon at the global level disables any individually configured and enabled sysmon tasks as well. However, any sysmon configuration that is made, including global and event-specific configuration are retained.

Examples The following example disables system monitoring:

```
Brocade# disable system-monitoring all
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands enable system-monitoring all

enable system-monitoring all

Enables system monitoring at the global level for all event types.

Syntax enable system-monitoring all

Command Default By default, system monitoring is enabled.

Command Modes Privileged EXEC mode.

Usage Guidelines This command enables system monitoring globally, and covers all event-specific system monitoring configuration as well. If specific configuration is not made for different types, default values defined at the global level are used.

Examples The following example enables all system monitoring tasks at the global level:

```
Brocade# enable system-monitoring all
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands disable system-monitoring all

sysmon timer

Configures the global system monitoring timer.

Syntax sysmon timer *minutes*

Command Default By default, the system monitoring timer is set to three minutes.

Parameters

minutes Specifies the system monitoring timer in minutes. The range of values is 1 through 60. The default value is 3.

Command Modes Global configuration mode.

Examples The following example sets the system monitoring timer to five minutes:

```
Brocade(config)# sysmon timer 5
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands sysmon threshold, sysmon log-backoff

sysmon log-backoff

Defines the number of times to skip logging an event before logging again at the global level. The **no** form of this command resets the parameter to default value.

Syntax **sysmon log-backoff** *number*

no sysmon log-backoff

Parameters *number* Specifies the number of times to skip an event logging before logging again.

Command Modes Global configuration mode.

Usage Guidelines Logging every error may not provide any new information, but adds significantly to the number of error entries that need to be analyzed. You can configure the system monitoring utility to ignore a certain number of errors (within a stream of consecutive errors) before writing the entry to the log again.

This option helps you further isolate issues that randomly occur from issues because of device failure. The sysmon utility keeps a counter of the number of times the threshold value is exceed. If the number exceeds the back-off value, the error is logged as specified by the action option.

Examples The following example sets the number of times to skip logging to 20.

```
Brocade(config)# sysmon log-backoff 20
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands sysmon threshold

sysmon threshold

Defines the threshold for errors at the global level. The **no** form of this command resets the threshold configuration to default values.

Syntax **sysmon threshold** *events polling-interval*

no sysmon threshold

Parameters

events	Specifies the threshold in terms of the number of events. Valid values are 1 through 10. When expressed in the command, the default value is 2.
polling-interval	Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events. When expressed in the command, the default value is 10.

Command Modes Global configuration mode.

Usage Guidelines The type-specific threshold values that you define overrides the global threshold value for each event. However, if you define the global value later, the latest value prevails. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and the action will always be triggered.

Examples The following example sets the threshold to 3 events over 7 consecutive polling periods:

```
Brocade(config)# sysmon threshold 3 7
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands sysmon fa error-count, sysmon fa link, sysmon pp error-count, sysmon xbar error-count, sysmon xbar link,

System monitoring on FCX and ICX devices

On FCX and ICX devices, system monitoring monitors the following errors:

- ECC errors.
- Link errors.

These errors are monitored on a stack unit basis.

Use the following commands configure and display the status of system monitoring on fabric adaptors:

- sysmon ecc-error
- sysmon link-error

sysmon ecc-error

Configures how sysmon handles ECC errors. The **no** version of this command disables system monitoring on internal ECC errors.

Syntax **sysmon ecc-error-count {threshold** *events polling-interval* | **log-backoff** *value* | **action {none** | **syslog}}**

no sysmon fa error-count

Command Default By default, the command takes the global values defined for threshold, log back-off, and action.

Parameters

threshold	Defines the threshold for errors. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and the action will always be triggered.
events	Specifies the threshold in terms of the number of events. Valid values are 1 through 10.
polling-interval	Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.
log-backoff	If an error condition persists, it will be continuously logged (internally and/or externally to syslog as defined by the action). The log back-off count skips configured number of logs before logging again.
action	Specifies the action to take when error count exceeds the specified threshold and log back-off values.
none	The error is logged in the internal sysmon logs. This is the default value.
syslog	The error is logged to syslog.

Command Modes Global configuration mode.

Usage Guidelines This command is supported only on FCX and ICX devices.

Examples The following example configures system monitoring for fabric adaptor errors:

```
Brocade(config)# sysmon ecc-error threshold 3 7
Brocade(config)# sysmon ecc-error action syslog
Brocade(config)# sysmon ecc-error log-backoff 15
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands sysmon link-error, sysmon timer, sysmon threshold

sysmon link-error

Configures how sysmon handles link errors. The **no** version of this command disables system monitoring on link errors.

Syntax **sysmon link-error** {**threshold** *events polling-interval* | **log-backoff** *value* | **action** {**none** | **syslog**}}

no sysmon link-error

Command Default By default, the command takes the global values defined for threshold, log back-off, and action.

Parameters

threshold Defines the threshold for errors. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and the action will always be triggered.

events Specifies the threshold in terms of the number of events. Valid values are 1 through 10.

polling-interval Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.

log-backoff If an error condition persists, it will be continuously logged (internally and/or externally to syslog as defined by the action). The log back-off count skips configured number of logs before logging again.

action Specifies the action to take when the error count exceeds the specified threshold and log back-off values.

none The error is logged in the internal sysmon logs. This is the default value.

syslog The error is logged to syslog.

Command Modes Global configuration mode.

Usage Guidelines This command is supported only on FCX and ICX devices.

Examples The following example configures system monitoring for fabric adaptor errors:

```
Brocade(config)# sysmon link-error threshold 3 7
Brocade(config)# sysmon link-error action syslog
Brocade(config)# sysmon link-error log-backoff 15
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands sysmon ecc-error, sysmon timer, sysmon threshold

System monitoring for Fabric Adapters

On FSX devices, system monitoring for fabric adaptors monitor errors such as the following:

- End of Packet (EoP) or Start of Packet (SoP) errors
- Cyclic Redundancy Check (CRC) errors
- Packets dropped due to congestion

In addition to the error count, sysmon also checks for connectivity of FA links. This happens at the interval defined by the sysmon-timer command generally or specifically for FA.

Use the following commands configure and display the status of system monitoring on fabric adaptors:

- sysmon fa error-count
- sysmon fa link
- show sysmon counters
- show sysmon logs
- show sysmon config

sysmon fa error-count

Configures how sysmon handles fabric adaptor-related errors. The **no** version of this command disables system monitoring on fabric adaptors.

Syntax **sysmon fa error-count** {**threshold** *events polling-interval* | **log-backoff** *value* | **action** {**none** | **syslog**}}

no sysmon fa error-count

Command Default By default, the command takes the global values defined for threshold, log back-off, and action.

Parameters

threshold Defines the threshold for errors. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and the action will always be triggered.

events Specifies the threshold in terms of the number of events. Valid values are 1 through 10.

polling-interval Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.

log-backoff If an error condition persists, it will be continuously logged (internally and/or externally to syslog as defined by the action). The log back-off count skips configured number of logs before logging again.

action Specifies the action to take when a fabric adapter error count exceeds the specified threshold and log back-off values.

none The error is logged in the internal sysmon logs. This is the default value.

syslog The error is logged to syslog.

Command Modes Global configuration mode.

Usage Guidelines This command is supported only on FSX devices.

Examples The following example configures system monitoring for fabric adaptor errors:

```
Brocade(config)# sysmon fa error-count threshold 3 7
Brocade(config)# sysmon fa error-count action syslog
Brocade(config)# sysmon fa error-count log-backoff 15
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands: sysmon fa link, sysmon timer, sysmon threshold

sysmon fa link

Configures system monitoring for link errors on all or specified fabric adaptors. The **no** form of this command resets the parameters to default values.

Syntax **sysmon fa link** {**threshold** *events polling-interval* | **log-backoff** *value* | **action** {**none** | **syslog**}}

no sysmon fa link

Command Default By default, the command takes the global values defined for threshold, log back-off, and action.

Parameters

threshold	Defines the failure threshold for the fabric adapter link error event. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and no event will be triggered.
events	Specifies the threshold in terms of the number of events. Valid values are 1 through 10.
polling-interval	Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.
log-backoff	If an error condition persists, it will be continuously logged (internally and/or externally). The log back-off count skips configured number of logs before logging again. This avoids overflow of the internal log or of the syslog.
action	Specifies the action to take when a fabric adapter link error exceeds the specified threshold and log back-off values.
none	No action is taken. This is the default.
syslog	The error is logged to syslog.

Command Modes Global configuration mode.

Usage Guidelines This command is supported only on FSX devices.

Examples The following example configures the sysmon options for fabric adaptor links:

```
Brocade(config)# sysmon fa link threshold 3 7
Brocade(config)# sysmon fa link action syslog
Brocade(config)# sysmon fa link log-backoff 15
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands sysmon fa error-count, sysmon timer, sysmon threshold

System monitoring for Cross Bar

On FSX devices, errors typically detected in the cross bar include:

- Bad (IP) headers
- Bad length errors
- Reformat errors

Besides the error count, sysmon also checks for connectivity of SFM/XBAR links. This happens at the interval defined by the sysmon-timer command generally or specifically for cross bar.

Use the following commands to configure and display the statistics of cross bar or switch fabric module:

- sysmon xbar error-count
- sysmon xbar link
- show sysmon logs
- show sysmon counters
- show sysmon config
- show sysmon system sfm

sysmon xbar error-count

Configures system monitoring for cross bar errors. The **no** form of this command resets the parameters to default values.

Syntax **sysmon xbar error-count** {**threshold** *events polling-interval* | **log-backoff** *value* | **action** {**none** | **syslog**}}

no sysmon xbar error-count

Command Default By default, the command takes the global values defined for threshold, log back-off, and action.

Parameters

threshold	Defines the failure threshold for the cross bar error-count event. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and no event will be triggered.
events	Specifies the threshold in terms of the number of events. Valid values are 1 through 10.
polling-interval	Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.
log-backoff	If an error condition persists, it will be continuously logged (internally and/or externally). The log back-off count skips configured number of logs before logging again. This avoids overflow of the internal log or of the syslog.
action	Specifies the action to take when the error count exceeds the specified threshold and log back-off values.
none	No action is taken.
syslog	The error is logged to syslog.

Command Modes Global configuration mode.

Usage Guidelines This command is supported only on FSX devices.

Examples The following example configures system monitoring for cross bar errors.

```
Brocade(config)# sysmon xbar error-count threshold 3 7
Brocade(config)# sysmon xbar error-count action syslog
Brocade(config)# sysmon xbar error-count log-backoff 15
```

History

Release	Command History
8.0	This command was introduced.

Related Commands sysmon xbar link, sysmon timer, sysmon threshold

sysmon xbar link

Configures the sysmon parameters for the crossbar link. The **no** form of this command resets the parameters to default values.

Syntax **sysmon xbar link** {**threshold** *events polling-interval* |**log-backoff** *value* | **action** {**none** | **syslog**}}

no sysmon xbar link

Command Default By default, the command takes the global values defined for threshold, log back-off, and action.

Parameters

threshold	Defines the failure threshold for the fabric adapter error-count event. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and no event will be triggered.
events	Specifies the threshold in terms of the number of events. Valid values are 1 through 10.
polling-interval	Specifies the number of polling windows. The device polls the internal registers at the interval specified by the sysmon timer value. Valid values 1-32. However, the polling window number must be equal or greater than the number of events.
log-backoff	If an error condition persists, it will be continuously logged (internally and/or externally). The log back-off count skips configured number of logs before logging again. This avoids overflow of the internal log or of the syslog.
action	Specifies the action to take when the error count exceeds the specified threshold and log back-off values.
none	No action is taken.
syslog	The error is logged to syslog.

Command Modes Global configuration mode.

Usage Guidelines This command is supported only on FSX devices.

Examples The following example configures system monitoring for cross bar link errors:

```
Brocade(config)# sysmon xbar link threshold 3 7
Brocade(config)# sysmon xbar link action syslog
Brocade(config)# sysmon xbar link log-backoff 15
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands sysmon xbar error-count, sysmon timer, sysmon threshold

System monitoring for Packet Processors

On FSX devices, errors typically detected in packet processors include:

- Parity errors
- Error Checking Code (ECC) errors
- ConfigTable0 errors
- TCAM error
- TCAM action parity errors
- Token bucket priority parity errors
- State variable parity errors
- Link list RAM ECC errors
- FBUF RAM ECC errors
- Egress VLAN parity errors
- Ingress VLAN parity errors
- Layer 2 port isolation parity errors
- Layer 3 port isolation parity errors
- VIDX parity errors

Besides the error count, sysmon also checks for connectivity of SFM/XBAR links. This happens at the interval defined by the sysmon-timer command generally or specifically for cross bar.

Use the following commands to configure and display the statistics of cross bar or switch fabric module:

- sysmon pp error-count
- show sysmon logs
- show sysmon counters
- show sysmon config

sysmon pp error-count

Configures the sysmon monitoring parameters for error events in packet processors. The **no** form of this command resets the parameters to default values.

Syntax **sysmon pp error-count** {**threshold** *events polling-interval* | **log-backoff** *value* | **action** {**none** | **syslog**}}

no sysmon pp error-count

Command Default By default, system monitoring on packet processors is enabled, with the global default values.

Parameters

threshold	Defines the failure threshold for the fabric adapter error-count event. The threshold is defined as N/W, where N is the number of events, and W is the number of consecutive polling periods. When the threshold is reached, actions configured for this event type will take place. Note that a threshold of 1/W implies that there is no threshold, and no event will be triggered.
log-backoff	If an error condition persists, it will be continuously logged (internally and/or externally). The log back-off count skips configured number of logs before logging again. This avoids overflow of the internal log or of the syslog.
action	Specifies the action to take when the error count exceeds the specified threshold and log back-off values.
none	No action is taken. This is the default action.
syslog	The error is logged to syslog.

Command Modes Global configuration mode.

Usage Guidelines This is a global configuration for all packet processors— you cannot configure sysmon parameters for individual packet processors. However, you can display the logs for individual packet processors by specifying the packet processor identifier.

This command is supported only on FSX devices.

Examples The following example configures system monitoring on packet processors:

```
Brocade(config)# sysmon pp error-count threshold 3 7
Brocade(config)# sysmon pp error-count action syslog
Brocade(config)# sysmon pp error-count log-backoff 15
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands sysmon timer, sysmon threshold

clear sysmon counters

Clears sysmon counters for all or specific event types.

Syntax

clear sysmon counters all

clear sysmon counters fa {**error** | **link**} {**all** | *decimal*}

clear sysmon counters pp error {**all** | *decimal*}

clear sysmon counters xbar {**error** | **link**} {**all** | *decimal*}

clear sysmon counters {**ecc-error** | **link-error**}

Parameters

all	Clears all sysmon counters.
fa	Clears the fabric adaptor sysmon counters.
error	Clears the fabric adaptor error counters. You can specify all or a fabric adaptor, identified by the index.
link	Clears the fabric adaptor sysmon counters for links. You can specify all or a fabric adaptor identified by the index.
pp error	Clears packet processor sysmon counters. You can specify all or a packet processor identified by the index.
xbar	Clears cross bar sysmon counters for cross bar. You can specify all or a cross bar identified by the index.
error	Clears the cross bar sysmon error counters. You can specify all or a cross bar identified by the index.
link	Clears the cross bar sysmon counters for links. You can specify all or a cross bar identified by the index.
ecc-error	Clears the ECC error count on FCX and ICX devices. This option is not supported on FSX devices.
stack-unit	Specifies the stack unit on which errors to be cleared.
all	Specifies that all stack units are cleared of errors.
link-error	Clears the link error count on FCX and ICX devices. This option is not supported on FSX devices.
stack-unit	Specifies the stack unit on which errors to be cleared.
all	Specifies that all stack units are cleared of errors.

Command Modes

Global configuration mode.

Usage Guidelines

None.

Examples

The following example clears the fabric adaptor sysmon counters.

```
Brocade(config)# clear sysmon counters fa error all
```

clear sysmon counters

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands

show sysmon counters, show sysmon logs, disable system-monitoring all

show sysmon logs

Displays the entries written to syslog for all event types if the action specified is to log them into syslog. If the action specified is **none**, the sysmon logs display nothing.

Syntax **show sysmon logs**

Command Default All syslog messages are displayed.

Command Modes Privileged EXEC mode.

Global configuration mode.

Examples The following example displays the syslog entries that were made by sysmon if the action specified either at the global level or type level was to log the events to syslog. If the action specified was **none**, no syslog entries exist.

```
Brocade(config)# show sysmon logs
Aug  3 03:59:22:C:Sysmon:XBAR LINK: SFM1/XBAR1/FPORT0 -- NO SYNC
Aug  3 03:59:22:C:Sysmon:FA Link: SLOT9/FA16/Link0 -- HG.Link error
Aug  3 03:58:22:W:Sysmon:PP ERROR: SLOT4/PP6  error occurred
Aug  3 03:59:34:W:Sysmon:FA ERROR: SLOT1/FA0  error occurred
Aug  3 03:60:34:W:Sysmon:XBAR ERROR: SFM1/XBAR1/FPORT2 -error  occurred
```

The following table describes the output of this command:

TABLE 37 show sysmon log s command output fields

Field	Description
Date and time	Aug 3 03:59:22
Critical or Warning	A ‘C’ indicates a critical error and a ‘W’ indicates a warning.
Sysmon	Message coming from Sysmon
Event type	Possible values are FA ERROR, FA Link, XBAR ERROR, XBAR LINK, or PP ERROR
Component identifier	Identifies the component of the system where the error was detected
Error	A brief description of the error

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands show sysmon counters, clear sysmon counters

show sysmon counters

Displays sysmon counters for all or specific event types.

Syntax **show sysmon counters** *type* {**error** | **link**}

show sysmon counters {**ecc-error** | **link-error**}

Command Default All counters are displayed.

Parameters

type The event type for which sysmon counters are displayed. For FSX devices, the options are **all**, **fa** (fabric adapter), **pp** (packet processor), and **xbar** (cross bar). For FCX and ICX devices, the options are **ecc-error** and **link-error**. The default value is **all**.

error Displays the error counter for the specified event type.

link Displays the link error counters. You can specify either all or specific links.

ecc-error Displays the ECC error count on FCX and ICX devices. This option is not supported on FSX devices.

stack-unit Specifies the stack unit on which errors to be displayed.

all Displays errors for all stack units.

link-error Displays the link error count on FCX and ICX devices. This option is not supported on FSX devices.

stack-unit Specifies the stack unit on which errors to be displayed.

all Displays errors for all stack units.

Command Modes Privileged EXEC mode.

Global configuration mode.

Examples The following displays all fabric adaptor statistics on an FSX device:

```
Brocade# show sysmon counters fa link all
Sysmon FA HG.link error detected (number of times)
                  FA-link0          FA-link1      FA-link2      FA-link3
SLOT    FA-dev    Sync/FC(RX,TX)    Sync/FC(RX,TX)Sync/FC(RX,TX)Sync/FC(RX,TX)
1       0         0/(0,0)           0/(0,0)       0/(0,0)       0/(0,0)
2       2         0/(0,0)           0/(0,0)       0/(0,0)       0/(0,0)
9       16        1751/(1750,1750)  0/(0,0)       0/(0,0)       0/(0,0)
9       17        0/(0,0)           0/(0,0)       0/(0,0)       0/(0,0)
```

The following example displays the error events that sysmon has recorded for the fabric adaptor 0.

```
Brocade# show sysmon counters fa error 0
Sysmon error detected on: SLOT 1, FA 0(number of times)

****PUMA Device 0 VOQUnit0 error detect
Set 0 EnQ Drop detect = 0
Set 1 EnQ Drop detect = 0
Set 2 EnQ Drop detect = 0
Set 3 EnQ Drop detect = 0
```

```
tail drop detect = 0 filter drop detect = 0, ecc drop detect = 0

****PUMA Device 0 VOQUnit1 error detect
Set 0 EnQ Drop detect = 0
Set 1 EnQ Drop detect = 0
Set 2 EnQ Drop detect = 0
Set 3 EnQ Drop detect = 0
tail drop detect = 0 filter drop detect = 0, ecc drop detect = 0

****PUMA Device 0 CRX error detect
CRC detect = 0, Lost SOP.EOP detect = 0, no egress Buf detect = 0
fifo full detect = 0, UC congest detect = 0, MC congest detect = 0
bad buf alloc detect = 0, e2e drop detect = 0
```

The following example shows the crossbar errors for the switch fabric module 0.

```
Brocade# show sysmon counters xbar error 0
Sysmon SFM 1 xbar 0 HG.link Rx error detected (number of times)
HG.link   BadLen   BadHeader ReformatErr
   0       0          0          0
   1       0          0          0
   2       0          1          0
   3       0          0          0
   4       0          0          0
   5       0          0          0
   6       0          0          0
   7       0          0          0
   8       0          0          0
   9       0          0          0
   10      0          0          0
   11      0          0          0
```

The following example displays the cross bar link errors for the SFM module 0.

```
Brocade# show sysmon counters xbar link 0
Sysmon SFM 0 xbar 1 HG.link NO-SYNC detected (number of times)
HG.link   NO-SYNC
0          0
1          0
2          0
3          0
4          0
5          1757
6          0
7          0
8          0
9          0
10         0
11         0
```

The following example displays the error counter for the specified packet processor 0.

```
Brocade# show sysmon counter pp error 0
Sysmon error detected on: SLOT 1, PP 0(number of times)

****PUMA Device 0 Buffer SRAM error detect
Ingress buffer error detect = 0
Egress buffer error detect  = 1
```

```
****PUMA Device 0 Control SRAM error detect
CSU :  Parity error detect  = 0, ECC error detect  = 0
LPM0:  Parity error detect  = 0, ECC error detect  = 0
LPM1:  Parity error detect  = 0, ECC error detect  = 0
LPM2:  Parity error detect  = 0, ECC error detect  = 0
LPM3:  Parity error detect  = 0, ECC error detect  = 0
```

The following example displays all error counter data on an FCX device:

```
Brocade(config)#show sysmon counters all
Sysmon error detected on: Stacking Unit 1 (number of times)

****Stacking unit 1 (FCX) Link error detect
Port 24
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 25
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 26
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 27
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
==========================
Sysmon error detected on: Stacking Unit 2 (number of times)

****Stacking unit 2 (FCX) Link error detect
Port 24
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 25
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 26
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 27
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
==========================
Sysmon error detected on: Stacking Unit 3 (number of times)

****Stacking unit 3 (FCX) Link error detect
Port 24
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 25
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 26
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 27
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
==========================
Sysmon error detected on: Stacking Unit 4 (number of times)

****Stacking unit 4 (FCX) Link error detect
Port 24
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 25
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 26
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 27
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
==========================
```

```
Sysmon error detected on: Stacking Unit 5 (number of times)

****Stacking unit 5 (FCX) Link error detect
Port 24
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 25
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 26
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
Port 27
 Link error detect = 0 remote fault detect = 0 lane error detect = 0
==========================
Sysmon ECC error detected on: Stacking Unit 1 (number of times)

****Stacking unit 1 (ICX) ecc error detect
ECC one-time error detect = 0 ECC two-time error detect = 0
==========================
Sysmon ECC error detected on: Stacking Unit 2 (number of times)

****Stacking unit 2 (ICX) ecc error detect
ECC one-time error detect = 0 ECC two-time error detect = 0
==========================
Sysmon ECC error detected on: Stacking Unit 3 (number of times)

****Stacking unit 3 (ICX) ecc error detect
ECC one-time error detect = 0 ECC two-time error detect = 0
==========================
Sysmon ECC error detected on: Stacking Unit 4 (number of times)

****Stacking unit 4 (ICX) ecc error detect
ECC one-time error detect = 0 ECC two-time error detect = 0
==========================
Sysmon ECC error detected on: Stacking Unit 5 (number of times)

****Stacking unit 5 (ICX) ecc error detect
ECC one-time error detect = 0 ECC two-time error detect = 0
==========================
```

History

Release	**Command History**
08.0.00a	This command was introduced.

Related Commands

clear sysmon counters, show sysmon logs

show sysmon config

Displays the complete sysmon configuration, including the global configuration and the event-specific configuration.

Syntax show sysmon config

Command Modes User EXEC mode.

Privileged EXEC mode.

Examples The following command displays the sysmon configuration an FSX device. The global configuration is displayed first, followed by the configuration for specific events.

```
Brocade> show sysmon config
======================================
System Monitoring (Sysmon) is: enabled
Sysmon timer = 3 minutes
======================================
Threshold: Times error detected / Consecutive times event polling.
Log Backoff Number: Number of times skip log before log again.
======================================
Sysmon Event: FA_ERROR_COUNT (Enabled)
Threshold:          2/10
Log Backoff Number: 10
Action: log(internal) /syslog

Sysmon Event: FA_LINK (Enabled)
Threshold:          2/10
Log Backoff Number: 10
Action: log(internal) /syslog

Sysmon Event: XBAR_ERROR_COUNT (Enabled)
Threshold:          2/10
Log Backoff Number: 10
Action: log(internal) /syslog

Sysmon Event: XBAR_LINK (Enabled)
Threshold:          2/10
Log Backoff Number: 10
Action: log(internal) /syslog

Sysmon Event: PP_ERROR_COUNT (Enabled)
Threshold:          2/10
Log Backoff Number: 10
Action: log(internal) /syslog
```

The following example displays the sysmon configuration on an FCX device:

```
Brocade(config)#show sysmon config
======================================
System Monitoring (Sysmon) is: enabled
Sysmon timer = 3 minutes
======================================
Threshold: Times error detected / Consecutive times event polling.
Log Backoff Numner: Number of times skip log before log again.
======================================
```

```
Sysmon Event: LINK_STATUS (Enabled)
Threshold:          2/10
Log Backoff Number: 10
Action: log(internal) /syslog

Sysmon Event: ECC_STATS (Enabled)
Threshold:          2/10
Log Backoff Number: 10
Action: log(internal) /syslog
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands

show sysmon counters, show sysmon logs

show sysmon system sfm

Displays the status of the switch fabric modules.

Syntax **show sysmon system sfm** {**all** | *number*}

Parameters

all	Displays the statistics for all SFMs on the device.
number	Specifies the SFM ID for which the statistics is to be displayed.

Command Modes

User EXEC mode.

Privileged EXEC mode.

Global configuration mode.

Usage Guidelines This command is supported only on FSX devices.

Examples The following command displays the statistics for all SFMs on the device.

```
Brocade(config)# show sysmon system sfm all
SFM= 1,Xbar= 2
 X-link  Status  FlowCtrl FA-dev/Link  Status  FlowCtrl
     2      OK       0x0       19/0       OK       0x0
     3      OK       0x0       13/0       OK       0x0
     4      OK       0x0        0/1       OK        --
     5      OK       0x0        3/0       OK       0x0
     7      OK       0x0       10/1       OK        --
     8      OK       0x0        7/0       OK       0x0
     9      OK       0x0       17/0       OK       0x0
=======================================================
SFM= 1,Xbar= 3
 X-link  Status  FlowCtrl FA-dev/Link  Status  FlowCtrl
     1      OK       0x0       17/1       OK       0x0
     2      OK       0x0        3/1       OK       0x0
     4      OK       0x0        0/2       OK        --
     5      OK       0x0       19/1       OK       0x0
     7      OK       0x0       10/2       OK        --
    10      OK       0x0        7/1       OK       0x0
    11      OK       0x0       13/1       OK       0x0
=======================================================
```

History

Release	Command History
08.0.00a	This command was introduced.

Related Commands clear sysmon counters

Appendix

Syslog messages

A

Table 1 lists all of the Syslog messages. Note that some of the messages apply only to Layer 3 switches.

NOTE
This chapter does not list Syslog messages that can be displayed when a debug option is enabled.

The messages are listed by message level, in the following order, then by message type:

- Emergencies (none)
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational
- Debugging

TABLE 1 Brocade Syslog messages

Message level	Message	Explanation
Alert	*num-modules* modules and 1 power supply, need more power supply!!	Indicates that the chassis needs more power supplies to run the modules in the chassis. The *num-modules* parameter indicates the number of modules in the chassis.
Alert	Fan *num, location*, failed	A fan has failed. The *num* is the fan number. The *location* describes where the failed fan is in the chassis.
Alert	MAC Authentication failed for *mac-address* on *portnum*	RADIUS authentication was successful for the specified *mac-address* on the specified *portnum*; however, the VLAN returned in the RADIUS Access-Accept message did not refer to a valid VLAN or VLAN ID on the Brocade device. This is treated as an authentication failure.
Alert	MAC Authentication failed for *mac-address* on *portnum* (Invalid User)	RADIUS authentication failed for the specified *mac-address* on the specified *portnum* because the MAC address sent to the RADIUS server was not found in the RADIUS server users database.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Alert	MAC Authentication failed for *mac-address* on *portnum* (No VLAN Info received from RADIUS server)	RADIUS authentication was successful for the specified *mac-address* on the specified *portnum*; however, dynamic VLAN assignment was enabled for the port, but the RADIUS Access-Accept message did not include VLAN information. This is treated as an authentication failure.
Alert	MAC Authentication failed for *mac-address* on *portnum* (Port is already in another radius given vlan)	RADIUS authentication was successful for the specified *mac-address* on the specified *portnum*; however, the RADIUS Access-Accept message specified a VLAN ID, although the port had previously been moved to a different RADIUS-assigned VLAN. This is treated as an authentication failure.
Alert	MAC Authentication failed for *mac-address* on *portnum* (RADIUS given vlan does not exist)	RADIUS authentication was successful for the specified *mac-address* on the specified *portnum*; however, the RADIUS Access-Accept message specified a VLAN that does not exist in the Brocade configuration. This is treated as an authentication failure.
Alert	MAC Authentication failed for *mac-address* on *portnum* (RADIUS given VLAN does not match with TAGGED vlan)	Multi-device port authentication failed for the *mac-address* on a tagged port because the packet with this MAC address as the source was tagged with a VLAN ID different from the RADIUS-supplied VLAN ID.
Alert	Management module at slot *slot-num* state changed from *module-state* to *module-state*.	Indicates a state change in a management module. The *slot-num* indicates the chassis slot containing the module. The *module-state* can be one of the following: • active • standby • crashed • coming-up • unknown
Alert	OSPF LSA Overflow, LSA Type = *lsa-type*	Indicates an LSA database overflow. The *lsa-type* parameter indicates the type of LSA that experienced the overflow condition. The LSA type is one of the following: • 1 – Router • 2 – Network • 3 – Summary • 4 – Summary • 5 – External
Alert	OSPF Memory Overflow	OSPF has run out of memory.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Alert	System: Module in slot *slot-num* encountered PCI config read error: Bus *PCI-bus-number*, Dev *PCI-device-number*, Reg Offset *PCI-config-register-offset*.	The module encountered a hardware configuration read error.
Alert	System: Module in slot *slot-num* encountered PCI config write error: Bus *PCI-bus-number*, Dev *PCI-device-number*, Reg Offset *PCI-config-register-offset*.	The module encountered a hardware configuration write error.
Alert	System: Module in slot *slot-num* encountered PCI memory read error: Mem Addr *memory-address*	The module encountered a hardware memory read error. The *memory-address* is in hexadecimal format.
Alert	System: Module in slot *slot-num* encountered PCI memory write error: Mem Addr *memory-address*.	The module encountered a hardware memory write error. The *memory-address* is in hexadecimal format.
Alert	System: Module in slot *slot-num* encountered unrecoverable PCI bridge validation failure. Module will be deleted.	The module encountered an unrecoverable (hardware) bridge validation failure. The module will be disabled or powered down.
Alert	System: Module in slot *slot-num* encountered unrecoverable PCI config read failure. Module will be deleted.	The module encountered an unrecoverable hardware configuration read failure. The module will be disabled or powered down.
Alert	System: Module in slot *slot-num* encountered unrecoverable PCI config write failure. Module will be deleted.	The module encountered an unrecoverable hardware configuration write failure. The module will be disabled or powered down.
Alert	System: Module in slot *slot-num* encountered unrecoverable PCI device validation failure. Module will be deleted.	The module encountered an unrecoverable (hardware) device validation failure. The module will be disabled or powered down.
Alert	System: Module in slot *slot-num* encountered unrecoverable PCI memory read failure. Module will be deleted.	The module encountered an unrecoverable hardware memory read failure. The module will be disabled or powered down.
Alert	System: Module in slot *slot-num* encountered unrecoverable PCI memory write failure. Module will be deleted.	The module encountered an unrecoverable hardware memory write failure. The module will be disabled or powered down.
Alert	System: No Free Tcam Entry available. System will be unstable	You must reboot the device.
Alert	System: Temperature is over shutdown level, system is going to be reset in *num* seconds	The chassis temperature has risen above shutdown level. The system will be shut down in the amount of time indicated.
Alert	Temperature *degrees* C degrees, warning level *warn-degrees* C degrees, shutdown level *shutdown-degrees* C degrees	Indicates an over temperature condition on the active module. The *degrees* value indicates the temperature of the module. The *warn-degrees* value is the warning threshold temperature configured for the module. The *shutdown-degrees* value is the shutdown temperature configured for the module.

TABLE 1 Brocade **Syslog messages** (Continued)

Message level	Message	Explanation
Critical	Authentication shut down *portnum* due to DOS attack	Denial of Service (DoS) attack protection was enabled for multi-device port authentication on the specified *portnum*, and the per-second rate of RADIUS authentication attempts for the port exceeded the configured limit. The Brocade device considers this to be a DoS attack and disables the port.
Debug	BGP4: Not enough memory available to run BGP4	The device could not start the BGP4 routing protocol because there is not enough memory available.
Debug	DOT1X: Not enough memory	There is not enough system memory for 802.1X authentication to take place. Contact Brocade Technical Support.
Error	No of prefixes received from BGP peer *ip-addr* exceeds maximum prefix-limit...shutdown	The Layer 3 switch has received more than the specified maximum number of prefixes from the neighbor, and the Layer 3 switch is therefore shutting down its BGP4 session with the neighbor.
Informational	IPv6: IPv6 protocol disabled on the device from *session-id*	IPv6 protocol was disabled on the device during the specified session.
Informational	IPv6: IPv6 protocol enabled on the device from *session-id*	IPv6 protocol was enabled on the device during the specified session.
Informational	MAC Filter applied to port *port-id* by *username* from *session-id* (filter id=*filter-ids*)	Indicates a MAC address filter was applied to the specified port by the specified user during the specified session. *session-id* can be console, telnet, ssh, or snmp. *filter-ids* is a list of the MAC address filters that were applied.
Informational	MAC Filter removed from port *port-id* by *username* from *session-id* (filter id=*filter-ids*)	Indicates a MAC address filter was removed from the specified port by the specified user during the specified session. *session-id* can be console, telnet, ssh, or snmp. *filter-ids* is a list of the MAC address filters that were removed.
Informational	Security: Password has been changed for user *username* from *session-id*	Password of the specified user has been changed during the specified session ID or type. *session-id* can be console, telnet, ssh, or snmp.
Informational	*device-name*: Logical link on interface ethernet *slot#/port#* is down.	The specified ports were logically brought down while **singleton** was configured on the port.
Informational	*device-name*: Logical link on interface ethernet *slot#/port#* is up.	The specified ports were logically brought up while **singleton** was configured on the port.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Informational	*user-name* login to PRIVILEGED mode	A user has logged into the Privileged EXEC mode of the CLI. The *user-name* is the user name.
Informational	*user-name* login to USER EXEC mode	A user has logged into the USER EXEC mode of the CLI. The *user-name* is the user name.
Informational	*user-name* logout from PRIVILEGED mode	A user has logged out of Privileged EXEC mode of the CLI. The *user-name* is the user name.
Informational	*user-name* logout from USER EXEC mode	A user has logged out of the USER EXEC mode of the CLI. The *user-name* is the user name.
Informational	ACL *ACL id* added | deleted | modified from console | telnet | ssh| snmp session	A user created, modified, deleted, or applied an ACL through an SNMP, console, SSH, or Telnet session.
Informational	Bridge is new root, vlan *vlan-id*, root ID *root-id*	A Spanning Tree Protocol (STP) topology change has occurred, resulting in the Brocade device becoming the root bridge. The *vlan-id* is the ID of the VLAN in which the STP topology change occurred. The *root-id* is the STP bridge root ID.
Informational	Bridge root changed, vlan *vlan-id*, new root ID *string*, root interface *portnum*	A Spanning Tree Protocol (STP) topology change has occurred. The *vlan-id* is the ID of the VLAN in which the STP topology change occurred. The *root-id* is the STP bridge root ID. The *portnum* is the number of the port connected to the new root bridge.
Informational	Bridge topology change, vlan *vlan-id*, interface *portnum*, changed state to *stp-state*	A Spanning Tree Protocol (STP) topology change has occurred on a port. The *vlan-id* is the ID of the VLAN in which the STP topology change occurred. The *portnum* is the port number. The *stp-state* is the new STP state and can be one of the following: • disabled • blocking • listening • learning • forwarding • unknown
Informational	Cold start	The device has been powered on.
Informational	DHCP: snooping on untrusted port *portnum*, type number, drop	The device has indicated that the DHCP client receives DHCP server reply packets on untrusted ports, and packets are dropped.
Informational	DOT1X: port *portnum* - MAC *mac address* Cannot apply an ACL or MAC filter on a port member of a VE (virtual interface)	The RADIUS server returned an IP ACL or MAC address filter, but the port is a member of a virtual interface (VE).

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Informational	DOT1X: port *portnum* - MAC *mac address* cannot remove inbound ACL	An error occurred while removing the inbound ACL.
Informational	DOT1X: port *portnum* - MAC *mac address* Downloading a MAC filter, but MAC filter have no effect on router port	The RADIUS server returned an MAC address filter, but the *portnum* is a router port (it has one or more IP addresses).
Informational	DOT1X: port *portnum* - MAC *mac address* Downloading an IP ACL, but IP ACL have no effect on a switch port	The RADIUS server returned an IP ACL, but the *portnum* is a switch port (no IP address).
Informational	DOT1X:port *portnum* - MAC *mac address* Error - could not add all MAC filters	The Brocade device was unable to implement the MAC address filters returned by the RADIUS server.
Informational	DOT1X: port *portnum* - MAC *mac address* Invalid MAC filter ID - this ID doesn't exist	The MAC address filter ID returned by the RADIUS server does not exist in the Brocade configuration.
Informational	DOT1X: port *portnum* - MAC *mac address* Invalid MAC filter ID - this ID is user defined and cannot be used	The port was assigned a MAC address filter ID that had been dynamically created by another user.
Informational	DOT1X: port *portnum* - MAC *mac address* is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC address filters	802.1X authentication failed for the Client with the specified *mac address* on the specified *portnum* either due to insufficient system resources on the device, or due to invalid IP ACL or MAC address filter information returned by the RADIUS server.
Informational	DOT1X: port *portnum* - MAC *mac address* Port is already bound with MAC filter	The RADIUS server returned a MAC address filter, but a MAC address filter had already been applied to the port.
Informational	DOT1X:port *portnum* - MAC *mac address* This device doesn't support ACL with MAC Filtering on the same port	The RADIUS server returned a MAC address filter while an IP ACL was applied to the port, or returned an IP ACL while a MAC address filter was applied to the port.
Informational	DOT1X: Port *portnum* is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC address filters	802.1X authentication could not take place on the port. This happened because strict security mode was enabled and one of the following occurred: • Insufficient system resources were available on the device to apply an IP ACL or MAC address filter to the port • Invalid information was received from the RADIUS server (for example, the Filter-ID attribute did not refer to an existing IP ACL or MAC address filter)
Informational	DOT1X: Port *portnum* currently used vlan-id changes to *vlan-id* due to dot1x-RADIUS vlan assignment	A user has completed 802.1X authentication. The profile received from the RADIUS server specifies a VLAN ID for the user. The port to which the user is connected has been moved to the VLAN indicated by *vlan-id*.
Informational	DOT1X: Port *portnum* currently used vlan-id is set back to port default vlan-id *vlan-id*	The user connected to *portnum* has disconnected, causing the port to be moved back into its default VLAN, *vlan-id*.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Informational	DOT1X: Port *portnum*, AuthControlledPortStatus change: authorized	The status of the interface controlled port has changed from unauthorized to authorized.
Informational	DOT1X: Port *portnum*, AuthControlledPortStatus change: unauthorized	The status of the interface controlled port has changed from authorized to unauthorized.
Informational	Enable super | port-config | read-only password deleted | added | modified from console | telnet | ssh| snmp OR Line password deleted | added | modified from console | telnet | ssh| snmp	A user created, re-configured, or deleted an Enable or Line password through the SNMP, console, SSH, or Telnet session.
Informational	ERR_DISABLE: Interface ethernet *portnum* err-disable recovery timeout	Errdisable recovery timer expired and the port has been reenabled.
Informational	ERR_DISABLE: Interface ethernet 16, err-disable recovery timeout	If the wait time (port is down and is waiting to come up) expires and the port is brought up the following message is displayed.
Informational	ERR_DISABLE: Link flaps on port ethernet 16 exceeded threshold; port in err-disable state	The threshold for the number of times that a port link toggles from “up” to “down” and “down” to “up” has been exceeded.
Informational	Interface *portnum*, line protocol down	The line protocol on a port has gone down. The *portnum* is the port number.
Informational	Interface *portnum*, line protocol up	The line protocol on a port has come up. The *portnum* is the port number.
Informational	Interface *portnum*, state down	A port has gone down. The *portnum* is the port number.
Informational	Interface *portnum*, state up	A port has come up. The *portnum* is the port number.
Informational	MAC Based Vlan Disabled on port *port id*	A MAC Based VLAN has been disabled on a port
Informational	MAC Based Vlan Enabled on port *port id*	A MAC Based VLAN has been enabled on a port.
Informational	MAC Filter added | deleted | modified from console | telnet | ssh| snmp session filter id = *MAC filter ID*, src MAC = *Source MAC address* | any, dst MAC = *Destination MAC address* | any	A user created, modified, deleted, or applied this MAC address filter through the SNMP, console, SSH, or Telnet session.
Informational	MSTP: BPDU-guard interface ethernet *port-number* detect (Received BPDU), putting into err-disable state.	BPDU guard violation occurred in MSTP.
Informational	OPTICAL MONITORING: port *port-number* is not capable.	The optical transceiver is qualified by Brocade, but the transceiver does not support digital optical performance monitoring.
Informational	Port *p* priority changed to *n*	A port priority has changed.

TABLE 1 Brocade **Syslog messages** (Continued)

Message level	Message	Explanation
Informational	Port *portnum*, srcip-security max-ipaddr-per-int reached.Last IP=*ipaddr*	The address limit specified by the **srcip-security max-ipaddr-per-interface** command has been reached for the port.
Informational	Port *portnum*, srcip-security max-ipaddr-per-int reached.Last IP=*ipaddr*	The address limit specified by the **srcip-security max-ipaddr-per-interface** command has been reached for the port.
Informational	Security: console login by *username* to USER | PRIVILEGE EXEC mode	The specified user logged into the device console into the specified EXEC mode.
Informational	Security: console logout by *username*	The specified user logged out of the device console.
Informational	Security: telnet | SSH login by *username* from src IP *ip-address*, src MAC *mac-address* to USER | PRIVILEGE EXEC mode	The specified user logged into the device using Telnet or SSH from either or both the specified IP address and MAC address. The user logged into the specified EXEC mode.
Informational	Security: telnet | SSH logout by *username* from src IP *ip-address,* src MAC *mac-address* to USER | PRIVILEGE EXEC mode	The specified user logged out of the device. The user was using Telnet or SSH to access the device from either or both the specified IP address and MAC address. The user logged out of the specified EXEC mode.
Informational	SNMP read-only community | read-write community | contact | location | user | group | view | engineld | trap [host] [*value -str*] deleted | added | modified from console | telnet | ssh| snmp session	A user made SNMP configuration changes through the SNMP, console, SSH, or Telnet session. [*value-str*] does not appear in the message if SNMP **community** or **engineId** is specified.
Informational	SNMP Auth. failure, intruder IP: *ip-addr*	A user has tried to open a management session with the device using an invalid SNMP community string. The *ip-addr* is the IP address of the host that sent the invalid community string.
Informational	SSH | telnet server enabled | disabled from console | telnet | ssh| snmp session [by user *username*]	A user enabled or disabled an SSH or Telnet session, or changed the SSH enable/disable configuration through the SNMP, console, SSH, or Telnet session.
Informational	startup-config was changed or startup-config was changed by *user-name*	A configuration change was saved to the startup-config file. The *user-name* is the user ID, if they entered a user ID to log in.
Informational	STP: Root Guard Port *port-number,* VLAN *vlan-ID* consistent (Timeout).	Root guard unblocks a port.
Informational	STP: Root Guard Port *port-number*, VLAN *vlan-ID* inconsistent (Received superior BPDU).	Root guard blocked a port.
Informational	STP: VLAN *vlan id* BPDU-Guard on Port *port id* triggered (Received BPDU), putting into err-disable state	The BPDU guard feature has detected an incoming BPDU on {vlan-id, port-id}
Informational	STP: VLAN *vlan id* Root-Protect Port *port id*, Consistent (Timeout)	The root protect feature goes back to the consistent state.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Informational	STP: VLAN *vlan id* Root-Protect Port *port id*, Inconsistent (Received superior BPDU)	The root protect feature has detected a superior BPDU and goes into the inconsistent state on {*vlan-id*, *port-id*}.
Informational	STP: VLAN *vlan-id* BPDU-guard port *port-number* detect (Received BPDU), putting into err-disable state	STP placed a port into an errdisable state for BPDU guard.
Informational	STP: VLAN 1 BPDU-guard port *port-number* detect (Received BPDU), putting into err-disable state.	BPDU guard violation in occurred in STP or RSTP.
Informational	Syslog server *IP-address* deleted | added | modified from console | telnet | ssh| snmp OR Syslog operation enabled | disabled from console | telnet | ssh| snmp	A user made Syslog configuration changes to the specified Syslog server address, or enabled or disabled a Syslog operation through the SNMP, console, SSH, or Telnet session.
Informational	SYSTEM: Optic is not Brocade-qualified (*port-number*)	Brocade does not support the optical transceiver.
Informational	System: Fan *fan id* (from left when facing right side), ok	The fan status has changed from fail to normal.
Informational	System: Fan speed changed automatically to *fan speed*	The system automatically changed the fan speed to the speed specified in this message.
Informational	System: No free TCAM entry. System will be unstable	There are no TCAM entries available.
Informational	System: Static MAC entry with MAC Address *mac-address* is added from the *unit/slot/port* to *unit/slot/port* on VLANs *vlan-id* to *vlan-id*	A MAC address is added to a range of interfaces, which are members of the specified VLAN range.
Informational	System: Static MAC entry with MAC Address *mac-address* is added to the *unit/slot/port* to *unit/slot/port* on *vlan-id*	A MAC address is added to a range of interfaces, which are members of the specified VLAN.
Informational	System: Static MAC entry with MAC Address *mac-address* is added to portnumber *unit/slot/port* on VLAN *vlan-id*	A MAC address is added to an interface and the interface is a member of the specified VLAN.
Informational	System: Static MAC entry with MAC Address *mac-address* is deleted from the *unit/slot/port* to *unit/slot/port* on *vlan-id*	A MAC address is deleted from a range of interfaces, which are members of the specified VLAN.
Informational	System: Static MAC entry with MAC Address *mac-address* is deleted from et he *unit/slot/port* to *unit/slot/port* on VLANs *vlan-id* to *vlan-id*	A MAC address is deleted from a range of interfaces, which are members of the specified VLAN range.
Informational	System: Static MAC entry with MAC Address *mac-address* is deleted from portnumber *unit/slot/port* on *vlan-id*	A MAC address is deleted from an interface and the interface is a member of the specified VLAN.
Informational	System: Static MAC entry with MAC Address *mac-address* is deleted from portnumber *unit/slot/port* on VLANs *vlan-id* to *vlan-id*	A MAC address is deleted from an interface and the interface is a member of the specified VLAN range.

TABLE 1 Brocade **Syslog messages** (Continued)

Message level	Message	Explanation
Informational	telnet | SSH| access [by *username*] from src IP source *ip address*, src MAC *source MAC address* rejected, *n* attempts	There were failed SSH, or Telnet login access attempts from the specified source IP and MAC address. • [by user *username*] does not appear if **telnet** or **SSH** clients are specified. • *n* is the number of times this SNMP trap occurred in the last five minutes, or other configured number of minutes.
Informational	Trunk group (*ports*) created by 802.3ad link-aggregation module.	802.3ad link aggregation is configured on the device, and the feature has dynamically created a trunk group (aggregate link). The *ports* variable is a list of the ports that were aggregated to make the trunk group.
Informational	user *username* added | deleted | modified from console | telnet | ssh| snmp	A user created, modified, or deleted a local user account through the SNMP, console, SSH, or Telnet session.
Informational	vlan *vlan id* added | deleted | modified from console | telnet | ssh| snmp session	A user created, modified, or deleted a VLAN through the SNMP, console, SSH, or Telnet session.
Informational	Warm start	The system software (flash code) has been reloaded.
Informational	Stack: Stack unit *unit#* has been deleted to the stack system	The specified unit has been deleted from the stacking system.
Informational	Stack unit *unitNumber* has been elected as ACTIVE unit of the stack system	The specified unit in a stack has been elected as the Master unit for the stacking system.
Informational	Stack: Stack unit *unit#* has been added to the stack system	The specified unit has been added to the stacking system.
Informational	System: Management MAC address changed to *mac_address*	The management MAC address of a stacking system has been changed
Informational	System: Stack unit *unit#* Fan *fan#* (*description*), failed	The operational status of a fan in the specified unit in a stack changed from normal to failure.
Informational	System: Stack unit *unit#* Power supply *power-supply#* is down	The operational status of a power supply of the specified unit in a stack changed from normal to failure.
Informational	System: Stack unit *unit#* Power supply *power-supply#* is up	The operational status of a power supply of the specified unit in a stack changed from failure to normal.
Informational	System: Stack unit *unit#* Fan *fan#* (*description*), ok	The operational status of a fan in the specified unit in a stack changed from failure to normal.
Informational	System: Stack unit *unitNumber* Temperature *actual-temp* C degrees, warning level *warning-temp* C degrees, shutdown level *shutdown-temp* C degrees	The actual temperature reading for a unit in a stack is above the warning temperature threshold.

TABLE 1 Brocade **Syslog messages** (Continued)

Message level	Message	Explanation
Informational	vlan *vlan-id* Bridge is RootBridge *mac-address* (MgmtPriChg)	802.1W changed the current bridge to be the root bridge of the given topology due to administrative change in bridge priority.
Informational	vlan *vlan-id* Bridge is RootBridge *mac-address* (MsgAgeExpiry)	The message age expired on the Root port so 802.1W changed the current bridge to be the root bridge of the topology.
Informational	vlan *vlan-id* interface *portnum* Bridge TC Event (DOT1wTransition)	802.1W recognized a topology change event in the bridge. The topology change event is the forwarding action that started on a non-edge Designated port or Root port.
Informational	vlan *vlan-id* interface *portnum* STP state - *state* (DOT1wTransition)	802.1W changed the state of a port to a new state: forwarding, learning, blocking. If the port changes to blocking, the bridge port is in discarding state.
Informational	vlan *vlan-id* New RootBridge *mac-address* RootPort *portnum* (BpduRcvd)	802.1W selected a new root bridge as a result of the BPDUs received on a bridge port.
Informational	vlan *vlan-id* New RootPort *portnum* (RootSelection)	802.1W changed the port role to Root port, using the root selection computation.
Notification	ACL exceed max DMA L4 cam resource, using flow based ACL instead	The port does not have enough Layer 4 CAM entries for the ACL. To correct this condition, allocate more Layer 4 CAM entries. To allocate more Layer 4 CAM entries, enter the following command at the CLI configuration level for the interface: **ip access-group max-l4-cam** *num*
Notification	ACL insufficient L4 cam resource, using flow based ACL instead	The port does not have a large enough CAM partition for the ACLs
Notification	ACL insufficient L4 session resource, using flow based ACL instead	The device does not have enough Layer 4 session entries. To correct this condition, allocate more memory for sessions. To allocate more memory, enter the following command at the global CONFIG level of the CLI interface: **system-max session-limit** *num*
Notification	ACL port fragment packet inspect rate *rate* exceeded on port *portnum*	The fragment rate allowed on an individual interface has been exceeded. The *rate* indicates the maximum rate allowed. The *portnum* indicates the port. This message can occur if fragment thottling is enabled.
Notification	ACL system fragment packet inspect rate *rate* exceeded	The fragment rate allowed on the device has been exceeded. The *rate* indicates the maximum rate allowed. This message can occur if fragment thottling is enabled.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	Authentication Disabled on *portnum*	The multi-device port authentication feature was disabled on the on the specified *portnum*.
Notification	Authentication Enabled on *portnum*	The multi-device port authentication feature was enabled on the on the specified *portnum*.
Notification	BGP Peer *ip-addr* DOWN (IDLE)	Indicates that a BGP4 neighbor has gone down. The *ip-addr* is the IP address of the neighbor BGP4 interface with the Brocade device.
Notification	BGP Peer *ip-addr* UP (ESTABLISHED)	Indicates that a BGP4 neighbor has come up. The *ip-addr* is the IP address of the neighbor BGP4 interface with the Brocade device.
Notification	DHCP: snooping on untrusted port *portnum*, type number, drop	Indicates that the DHCP client receives DHCP server reply packets on untrusted ports, and packets are dropped.
Notification	DOT1X issues software but not physical port down indication of Port *portnum* to other software applications	The device has indicated that the specified is no longer authorized, but the actual port may still be active.
Notification	DOT1X issues software but not physical port up indication of Port *portnum* to other software applications	The device has indicated that the specified port has been authenticated, but the actual port may not be active.
Notification	DOT1X: Port *port_id* Mac *mac_address* -user *user_id* - RADIUS timeout for authentication	The RADIUS session has timed out for this 802.1x port.
Notification	ISIS L1 ADJACENCY DOWN *system-id* on circuit *circuit-id*	The Layer 3 switch adjacency with this Level-1 IS-IS has gone down. The *system-id* is the system ID of the IS-IS. The *circuit-id* is the ID of the circuit over which the adjacency was established.
Notification	ISIS L1 ADJACENCY UP *system-id* on circuit *circuit-id*	The Layer 3 switch adjacency with this Level-1 IS-IS has come up. The *system-id* is the system ID of the IS-IS. The *circuit-id* is the ID of the circuit over which the adjacency was established.
Notification	ISIS L2 ADJACENCY DOWN *system-id* on circuit *circuit-id*	The Layer 3 switch adjacency with this Level-2 IS-IS has gone down. The *system-id* is the system ID of the IS-IS. The *circuit-id* is the ID of the circuit over which the adjacency was established.
Notification	ISIS L2 ADJACENCY UP *system-id* on circuit *circuit-id*	The Layer 3 switch adjacency with this Level-2 IS-IS has come up. The *system-id* is the system ID of the IS-IS. The *circuit-id* is the ID of the circuit over which the adjacency was established.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	Local ICMP exceeds *burst-max* burst packets, stopping for *lockup* seconds!!	The number of ICMP packets exceeds the *burst-max* threshold set by the **ip icmp burst** command. The Brocade device may be the victim of a Denial of Service (DoS) attack. All ICMP packets will be dropped for the number of seconds specified by the *lockup* value. When the lockup period expires, the packet counter is reset and measurement is restarted.
Notification	Local TCP exceeds *burst-max* burst packets, stopping for *lockup* seconds!!	The number of TCP SYN packets exceeds the *burst-max* threshold set by the **ip tcp burst** command. The Brocade device may be the victim of a TCP SYN DoS attack. All TCP SYN packets will be dropped for the number of seconds specified by the *lockup* value. When the lockup period expires, the packet counter is reset and measurement is restarted.
Notification	Local TCP exceeds *num* burst packets, stopping for *num* seconds!!	Threshold parameters for local TCP traffic on the device have been configured, and the maximum burst size for TCP packets has been exceeded. The first *num* is the maximum burst size (maximum number of packets allowed). The second *num* is the number of seconds during which additional TCP packets will be blocked on the device. **NOTE:** This message can occur in response to an attempted TCP SYN attack.
Notification	MAC Authentication RADIUS timeout for *mac_address* on port *port_id*	The RADIUS session has timed out for the MAC address for this port.
Notification	MAC Authentication succeeded for *mac-address* on *portnum*	RADIUS authentication was successful for the specified *mac-address* on the specified *portnum*.
Notification	Module was inserted to slot *slot-num*	Indicates that a module was inserted into a chassis slot. The *slot-num* is the number of the chassis slot into which the module was inserted.
Notification	Module was removed from slot *slot-num*	Indicates that a module was removed from a chassis slot. The *slot-num* is the number of the chassis slot from which the module was removed.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF interface state changed, rid *router-id*, intf addr *ip-addr*, state *ospf-state*	Indicates that the state of an OSPF interface has changed. The *router-id* is the router ID of the Brocade device. The *ip-addr* is the interface IP address. The *ospf-state* indicates the state to which the interface has changed and can be one of the following: • down • loopback • waiting • point-to-point • designated router • backup designated router • other designated router • unknown
Notification	OSPF intf authen failure, rid *router-id*, intf addr *ip-addr*, pkt src addr *src-ip-addr*, error type *error-type*, pkt type *pkt-type*	Indicates that an OSPF interface authentication failure has occurred. The *router-id* is the router ID of the Brocade device. The *ip-addr* is the IP address of the interface on the Brocade device. The *src-ip-addr* is the IP address of the interface from which the Brocade device received the authentication failure. The *error-type* can be one of the following: • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown The *packet-type* can be one of the following: • hello • database description • link state request • link state update • link state ack • unknown

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF intf config error, *rid router-id*, intf addr *ip-addr*, pkt src addr *src-ip-addr*, error type *error-type*, pkt type *pkt-type*	Indicates that an OSPF interface configuration error has occurred. The *router-id* is the router ID of the Brocade device. The *ip-addr* is the IP address of the interface on the Brocade device. The *src-ip-addr* is the IP address of the interface from which the Brocade device received the error packet. The *error-type* can be one of the following: • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown The *packet-type* can be one of the following: • hello • database description • link state request • link state update • link state ack • unknown
Notification	OSPF intf rcvd bad pkt, rid *router-id*, intf addr *ip-addr*, pkt src addr *src-ip-addr*, pkt type *pkt-type*	Indicates that an OSPF interface received a bad packet. The *router-id* is the router ID of the Brocade device. The *ip-addr* is the IP address of the interface on the Brocade device. The *src-ip-addr* is the IP address of the interface from which the Brocade device received the authentication failure. The *packet-type* can be one of the following: • hello • database description • link state request • link state update • link state ack • unknown

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF intf rcvd bad pkt: Bad Checksum, rid *ip-addr*, intf addr *ip-addr*, pkt size *num*, checksum *num*, pkt src addr *ip-addr*, pkt type *type*	The device received an OSPF packet that had an invalid checksum. The rid *ip-addr* is the Brocade router ID. The intf addr *ip-addr* is the IP address of the Brocade interface that received the packet. The pkt size *num* is the number of bytes in the packet. The checksum *num* is the checksum value for the packet. The pkt src addr *ip-addr* is the IP address of the neighbor that sent the packet. The pkt type *type* is the OSPF packet type and can be one of the following: • hello • database description • link state request • link state update • link state acknowledgement • unknown (indicates an invalid packet type)
Notification	OSPF intf rcvd bad pkt: Bad Packet type, rid *ip-addr,* intf addr *ip-addr*, pkt size *num*, checksum *num*, pkt src addr *ip-addr*, pkt type *type*	The device received an OSPF packet with an invalid type. The parameters are the same as for the Bad Checksum message. The pkt type *type* value is “unknown”, indicating that the packet type is invalid.
Notification	OSPF intf rcvd bad pkt: Invalid packet size, rid *ip-addr,* intf addr *ip-addr,* pkt size *num*, checksum *num*, pkt src addr *ip-addr*, pkt type *type*	The device received an OSPF packet with an invalid packet size. The parameters are the same as for the Bad Checksum message.
Notification	OSPF intf rcvd bad pkt: Unable to find associated neighbor, rid *ip-addr,* intf addr *ip-addr,* pkt size *num*, checksum *num*, pkt src addr *ip-addr*, pkt type *type*	The neighbor IP address in the packet is not in the list of OSPF neighbors in the Brocade device. The parameters are the same as for the Bad Checksum message.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF intf retransmit, rid *router-id,* intf addr *ip-addr,* nbr rid nbr-*router-id*, pkt type is *pkt-type,* LSA type *lsa-type*, LSA id *lsa-id,* LSA rid *lsa-router-id*	An OSPF interface on the Brocade device has retransmitted a Link State Advertisement (LSA). The *router-id* is the router ID of the Brocade device. The *ip-addr* is the IP address of the interface on the Brocade device. The *nbr-router-id* is the router ID of the neighbor router. The *packet-type* can be one of the following: • hello • database description • link state request • link state update • link state ack • unknown The *lsa-type* is the type of LSA. The *lsa-id* is the LSA ID. The *lsa-router-id* is the LSA router ID.
Notification	OSPF LSDB approaching overflow, rid *router-id*, limit *num*	The software is close to an LSDB condition. The *router-id* is the router ID of the Brocade device. The *num* is the number of LSAs.
Notification	OSPF LSDB overflow, rid *router-id,* limit *num*	A Link State Database Overflow (LSDB) condition has occurred. The *router-id* is the router ID of the Brocade device. The *num* is the number of LSAs.
Notification	OSPF max age LSA, rid *router-id*, area *area-id*, LSA type *lsa-type*, LSA id *lsa-id*, LSA rid *lsa-router-id*	An LSA has reached its maximum age. The *router-id* is the router ID of the Brocade device. The *area-id* is the OSPF area. The *lsa-type* is the type of LSA. The *lsa-id* is the LSA ID. The *lsa-router-id* is the LSA router ID.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF nbr state changed, rid *router-id*, nbr addr *ip-addr*, nbr rid *nbr-router-Id*, state *ospf-state*	Indicates that the state of an OSPF neighbor has changed. The *router-id* is the router ID of the Brocade device. The *ip-addr* is the IP address of the neighbor. The *nbr-router-id* is the router ID of the neighbor. The *ospf-state* indicates the state to which the interface has changed and can be one of the following: • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown
Notification	OSPF originate LSA, rid *router-id*, area *area-id*, LSA type *lsa-type*, LSA id *lsa-id*, LSA router id *lsa-router-id*	An OSPF interface has originated an LSA. The *router-id* is the router ID of the Brocade device. The *area-id* is the OSPF area. The *lsa-type* is the type of LSA. The *lsa-id* is the LSA ID. The *lsa-router-id* is the LSA router ID.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF virtual intf authen failure, rid *router-id*, intf addr *ip-addr*, pkt src addr *src-ip-addr*, error type *error-type*, pkt type *pkt-type*	Indicates that an OSPF virtual routing interface authentication failure has occurred. The *router-id* is the router ID of the Brocade device. The *ip-addr* is the IP address of the interface on the Brocade device. The src-ip-addr is the IP address of the interface from which the Brocade device received the authentication failure. The *error-type* can be one of the following: • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown The *packet-type* can be one of the following: • hello • database description • link state request • link state update • link state ack • unknown

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF virtual intf config error, rid *router-id*, intf addr *ip-addr*, pkt src addr *src-ip-addr*, error type *error-type*, pkt type *pkt-type*	Indicates that an OSPF virtual routing interface configuration error has occurred. The *router-id* is the router ID of the Brocade device. The *ip-addr* is the IP address of the interface on the Brocade device. The *src-ip-addr* is the IP address of the interface from which the Brocade device received the error packet. The *error-type* can be one of the following: • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown The *packet-type* can be one of the following: • hello • database description • link state request • link state update • link state ack • unknown
Notification	OSPF virtual intf rcvd bad pkt, rid *router-id*, intf addr *ip-addr*, pkt src addr *src-ip-addr*, pkt type *pkt-type*	Indicates that an OSPF interface received a bad packet. The *router-id* is the router ID of the Brocade device. The *ip-addr* is the IP address of the interface on the Brocade device. The *src-ip-addr* is the IP address of the interface from which the Brocade device received the authentication failure. The *packet-type* can be one of the following: • hello • database description • link state request • link state update • link state ack • unknown

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF virtual intf retransmit, rid *router-id*, intf addr *ip-addr*, nbr rid *nbr-router-id*, pkt type is *pkt-type*, LSA type *lsa-type*, LSA id *lsa-id*, LSA rid *lsa-router-id*	An OSPF interface on the Brocade device has retransmitted a Link State Advertisement (LSA). The *router-id* is the router ID of the Brocade device. The *ip-addr* is the IP address of the interface on the Brocade device. The *nbr-router-id* is the router ID of the neighbor router. The *packet-type* can be one of the following: • hello • database description • link state request • link state update • link state ack • unknown The *lsa-type* is the type of LSA. The *lsa-id* is the LSA ID. The *lsa-router-id* is the LSA router ID.
Notification	OSPF virtual intf state changed, rid *router-id*, area *area-id*, nbr *ip-addr*, state *ospf-state*	Indicates that the state of an OSPF virtual routing interface has changed. The *router-id* is the router ID of the router the interface is on. The *area-id* is the area the interface is in. The *ip-addr* is the IP address of the OSPF neighbor. The *ospf-state* indicates the state to which the interface has changed and can be one of the following: • down • loopback • waiting • point-to-point • designated router • backup designated router • other designated router • unknown

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	OSPF virtual nbr state changed, rid *router-id*, nbr addr *ip-addr*, nbr rid *nbr-router-id*, state *ospf-state*	Indicates that the state of an OSPF virtual neighbor has changed. The *router-id* is the router ID of the Brocade device. The *ip-addr* is the IP address of the neighbor. The *nbr-router-id* is the router ID of the neighbor. The *ospf-state* indicates the state to which the interface has changed and can be one of the following: • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown
Notification	Transit ICMP in interface *portnum* exceeds *num* burst packets, stopping for *num* seconds!!	Threshold parameters for ICMP transit (through) traffic have been configured on an interface, and the maximum burst size for ICMP packets on the interface has been exceeded. The *portnum* is the port number. The first *num* is the maximum burst size (maximum number of packets allowed). The second *num* is the number of seconds during which additional ICMP packets will be blocked on the interface. **NOTE:** This message can occur in response to an attempted Smurf attack.
Notification	Transit TCP in interface *portnum* exceeds *num* burst packets, stopping for *num* seconds!	Threshold parameters for TCP transit (through) traffic have been configured on an interface, and the maximum burst size for TCP packets on the interface has been exceeded. The *portnum* is the port number. The first *num* is the maximum burst size (maximum number of packets allowed). The second *num* is the number of seconds during which additional TCP packets will be blocked on the interface. **NOTE:** This message can occur in response to an attempted TCP SYN attack.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Notification	VRRP intf state changed, intf *portnum*, vrid *virtual-router-id*, state *vrrp-state* VRRP (IPv6) intf state changed, intf *portnum*, vrid *virtual-router-id*, state *vrrp-state*	A state change has occurred in a Virtual Router Redundancy Protocol (VRRP) or VRRP-E IPv4 or IPv6 interface. The *portnum* is the port or interface where VRRP or VRRP-E is configured. The *virtual-router-id* is the virtual router ID (VRID) configured on the interface. The *vrrp-state* can be one of the following: • init • master • backup • unknown
Warning	DOT1X security violation at port *portnum*, malicious MAC address detected: *mac-address*	A security violation was encountered at the specified port number.
Warning	Dup IP *ip-addr* detected, sent from MAC *mac-addr* interface *portnum*	Indicates that the Brocade device received a packet from another device on the network with an IP address that is also configured on the Brocade device. The *ip-addr* is the duplicate IP address. The *mac-addr* is the MAC address of the device with the duplicate IP address. The *portnum* is the Brocade port that received the packet with the duplicate IP address. The address is the packet source IP address.
Warning	IGMP/MLD no hardware vidx, broadcast to the entire vlan. rated limited number	IGMP or MLD snooping has run out of hardware application VLANs. There are 4096 application VLANs per device. Traffic streams for snooping entries without an application VLAN are switched to the entire VLAN and to the CPU to be dropped. This message is rate-limited to appear a maximum of once every 10 minutes. The rate-limited number shows the number on non-printed warnings.
Warning	IGMP/MLD: *vlanId(portId)* is V1 but rcvd V2 from nbr *ipAddr*	Port has received a query with a MLD version that does not match the port MLD version. This message is rated-limited to appear a maximum of once every 10 hours.
Warning	Latched low RX Power | TX Power | TX Bias Current | Supply Voltage | Temperature warning alarm | warning, port *port-number*	The optical transceiver on the given port has risen above or fallen below the alarm or warning threshold.

TABLE 1 Brocade Syslog messages (Continued)

Message level	Message	Explanation
Warning	list *ACL-num* denied *ip-proto* *src-ip-addr* (*src-tcp/udp-port*) (Ethernet *portnum mac-addr*) - *dst-ip-addr* (*dst-tcp/udp-port*), 1 event(s)	Indicates that an Access Control List (ACL) denied (dropped) packets. The *ACL-num* indicates the ACL number. Numbers 1 – 99 indicate standard ACLs. Numbers 100 – 199 indicate extended ACLs. The *ip-proto* indicates the IP protocol of the denied packets. The *src-ip-addr* is the source IP address of the denied packets. The *src-tcp/udp-port* is the source TCP or UDP port, if applicable, of the denied packets. The *portnum* indicates the port number on which the packet was denied. The *mac-addr* indicates the source MAC address of the denied packets. The *dst-ip-addr* indicates the destination IP address of the denied packets. The *dst-tcp/udp-port* indicates the destination TCP or UDP port number, if applicable, of the denied packets.
Warning	Locked address violation at interface *eportnum*, address *mac-address*	Indicates that a port on which you have configured a lock-address filter received a packet that was dropped because the packet source MAC address did not match an address learned by the port before the lock took effect. The *eportnum* is the port number. The *mac-address* is the MAC address that was denied by the address lock. Assuming that you configured the port to learn only the addresses that have valid access to the port, this message indicates a security violation.
Warning	MAC filter group denied packets on port *portnum,* src macaddr *mac-addr*, *num* packets	Indicates that a MAC address filtergroup configured on a port has denied packets. The *portnum* is the port on which the packets were denied. The *mac-addr* is the source MAC address of the denied packets. The *num* indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.
Warning	multicast no software resource: *resource-name*, rate-limited *number*	IGMP or MLD snooping has run out of software resources. This message is rate-limited to appear a maximum of once every 10 minutes. The rate-limited *number* shows the number of non-printed warnings.

TABLE 1 Brocade **Syslog messages** (Continued)

Message level	Message	Explanation
Warning	No global IP! cannot send IGMP msg.	The device is configured for **ip multicast active** but there is no configured IP address and the device cannot send out IGMP queries.
Warning	No of prefixes received from BGP peer *ip-addr* exceeds warning limit *num*	The Layer 3 switch has received more than the allowed percentage of prefixes from the neighbor. The *ip-addr* is the IP address of the neighbor. The *num* is the number of prefixes that matches the percentage you specified. For example, if you specified a threshold of 100 prefixes and 75 percent as the warning threshold, this message is generated if the Layer 3 switch receives a 76th prefix from the neighbor.
Warning	rip filter list *list-num direction* V1 | V2 denied *ip-addr*, *num* packets	Indicates that a RIP route filter denied (dropped) packets. The *list-num* is the ID of the filter list. The *direction* indicates whether the filter was applied to incoming packets or outgoing packets. The value can be one of the following: • in • out The V1 or V2 value specifies the RIP version (RIPv1 or RIPv2). The *ip-addr* indicates the network number in the denied updates. The *num* indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.
Warning	Temperature is over warning level.	The chassis temperature has risen above the warning level.

Index

Numerics

A

B

C

D

E

F

H

I

L

M

N

O

P

R

S

T

V

X